Using Liferay DXP

A Complete Guide

THE LIFERAY DOCUMENTATION TEAM
Richard Sezov, Jr.
Jim Hinkey
Stephen Kostas
Jesse Rao
Cody Hoag
Nicholas Gaskill
Michael Williams

Liferay Press

Using Liferay DXP 7.1
by The Liferay Documentation Team

This book was created out of material from the Liferay Docs repository. Where the content of this book and the repository differ, the site is more up to date.

Contents

Preface

Welcome to the world of Liferay DXP! This book was written for anyone who has any part in setting up, using, or maintaining a web site built on Liferay DXP. For the end user, it contains everything you need to know about using the applications included with Liferay. For the administrator, you'll learn all you need to know about setting up your site with users, sites, organizations, and user groups, as well as how to manage your site's security with roles. For server admins, it guides you step-by-step through the installation, configuration, and optimization of Liferay DXP, including setting it up in a clustered, enterprise-ready environment. Use this book as a handbook for everything you need to do to get your Liferay DXP installation running smoothly, and then keep it by your side as you configure and maintain your Liferay-powered web site.

Conventions

The information contained herein has been organized in a way that makes it easy to locate information. The book has two parts. The first part, *Using Liferay Portal*, describes how to configure and use a freshly installed Liferay Portal. The second part, *Deploying Liferay Portal*, is for administrators who want to install Liferay Portal and optimize its performance.

Sections are broken up into multiple levels of headings, and these are designed to make it easy to find information.

Source code and configuration file directives are presented monospaced, as below.

```
Source code appears in a non-proportional font.
```

Italics represent links or buttons to be clicked on in a user interface.

`Monospaced type` denotes Java classes, code, or properties within the text.

Bold describes field labels and portlets.

Page headers denote the chapters and the section within the chapter.

Publisher Notes

It is our hope that this book is valuable to you, and that it becomes an indispensable resource as you work with Liferay DXP. If you need assistance beyond what is covered in this book, Liferay offers training[1], consulting[2], and support[3] services to fill any need that you might have.

For up-to-date documentation on the latest versions of Liferay, please see the documentation pages on Liferay Learn.[4]

As always, we welcome feedback. If there is any way you think we could make this book better, please feel free to mention it on our forums or in the feedback on Liferay Learn. You can also use any of the email addresses on our Contact Us page.[5] We are here to serve you, our users and customers, and to help make your experience using Liferay DXP the best it can be.

[1]https://learn.liferay.com
[2]https://www.liferay.com/consulting
[3]https://help.liferay.com
[4]https://learn.liferay.com
[5]https://www.liferay.com/contact-us

Part I

Using Liferay DXP

Chapter 1

What is Liferay DXP

Choosing a platform for your web presence can be daunting. You want something built on a solid technology foundation, fully featured, fully customizable, easy to use, with a solid record of success, and a forward-looking, bright vision for your site's future. If you're a small organization, the interface should be intuitive enough and the defaults safe enough for you to get your site up quickly, so you can focus on your mission instead of your technology. If you're a large organization, your marketers should be able to update your site's content easily, your developers should have a solid platform on which to build your applications, and your administrators should have all the tools necessary to manage the system and connect it to your existing environment.

You've just found all that in Liferay DXP.

Liferay DXP provides a robust *platform* to build your site on quickly and serve it to all clients—desktop, mobile, or anything in between. It provides all the standard *applications* you need your site to run. It provides an easy to use development *framework* for new applications or customization. On top of this, Liferay DXP is developed using an open source methodology, by people from around the world. The code base is solid, reliable, and stable in mission critical deployments in diverse industries.

The best way to explain all this is to show you how Liferay DXP works in four key areas:

- Web Content Management
- User collaboration
- Development Platform
- Customization

Of course, this is only part of what Liferay DXP can do, but it suffices as an introduction to how it works. Read on to learn about Liferay DXP's web content management features.

1.1 Building a site with Liferay Web Experience

When you log into Liferay DXP and look at its default screen, one thing to notice is that it's built for all clients that access the web, not just desktop browsers.

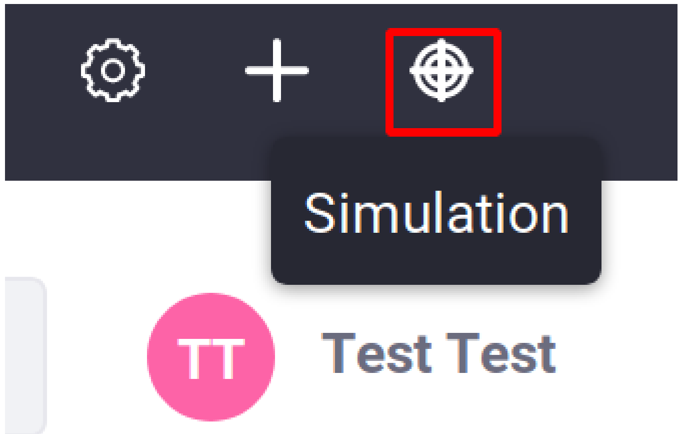

Figure 1.1: If you click this button, a preview appears that lets you see how the page would look when displayed at various device resolutions.

All the page controls are right there for you. For example, to add content or applications to the page, click the + button, and then you can add any kind of content Liferay DXP supports. The interface gets out of your way so you can do your work.

Liferay Web Experience Management scales to work for the tiniest of sites all the way up to the largest of sites. You can get started quickly by clicking that *Add* button and choosing *Add New* → *Basic Web Content*, where you can immediately start typing content into a WYSIWYG editor, in place.

If you are comfortable with developing websites, you can instead click the Menu button on the left (), select *Build* → *Page Fragments*, and use the Fragment Editor to build a page out of HTML, CSS, and JavaScript.

Large organizations can use Liferay DXP to host many different websites, all with their own domain names. Each site can take advantage of a separate staging server, where content and pages are created by teams of people using structures and templates, and updates to the production server are published on a schedule, only after having gone through a multi-step approval process. That's powerful.

By default, Liferay DXP starts with a single site that has a single page. You can build any website you wish out of this, complete with multi-nested page hierarchies, as the figure below shows.

Pages can be added, removed, or reordered any time, and you have the full flexibility of all the HTML page attributes like meta tags and robot file declarations, as well as full control over the layout.

Pages are also integrated with Liferay's powerful permissions system, so it's easy to restrict access to certain portions of your site. You can give individual users sites of their own, with public pages that have their content and blog and private pages that contain their calendars and profiles.

If you run a large website with lots of different sub-sites for individuals and groups, you can use page templates and site templates. Page templates have layouts and applications already on them, and site template predefine a whole site made up of multiple page templates.

Web content is one example of an *asset*. Assets are meta-data attached to content types, and that meta-data aggregates similar assets together in searches or as published content. One way to do this is to tag and categorize content so it can be found more easily by users.

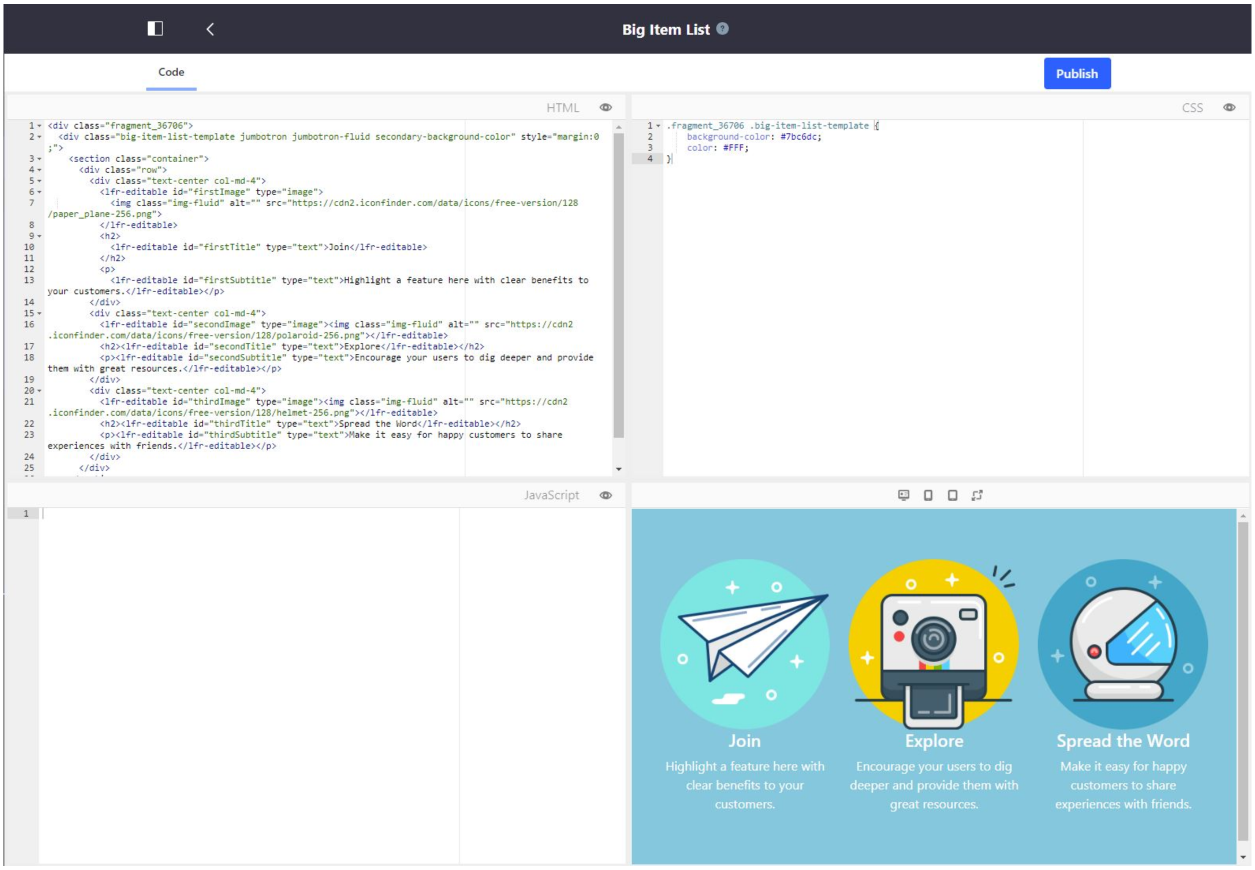

Figure 1.2: Fragments comprise HTML, CSS, and JavaScript, and can be used as building blocks for pages.

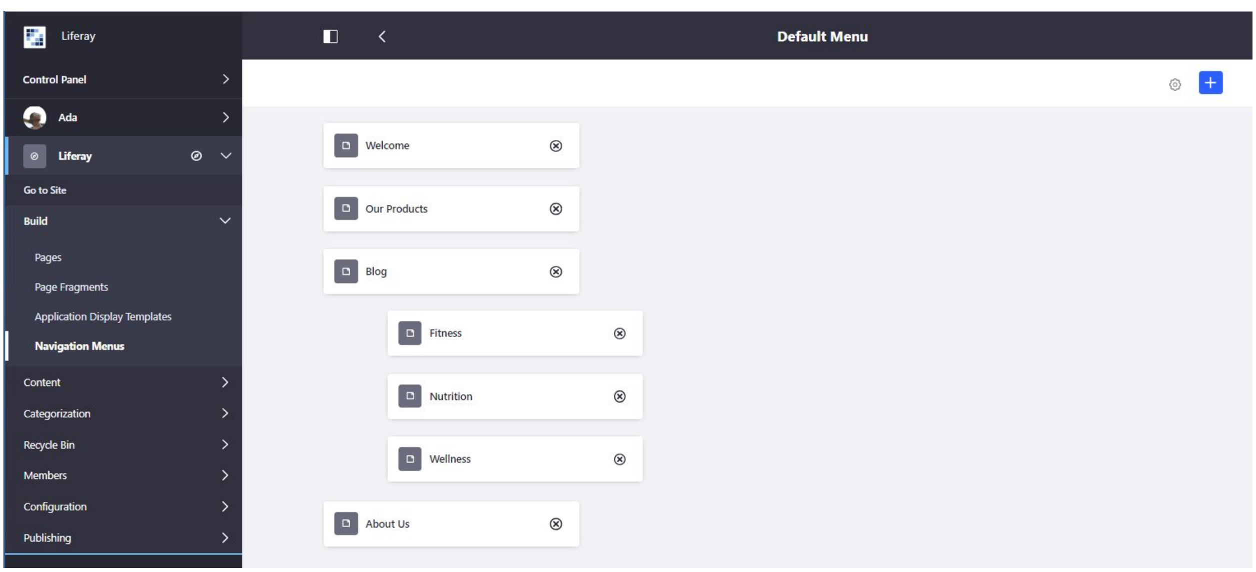

Figure 1.3: Liferay's page hierarchies are easy to create, using a tree structure that's familiar to anyone who has used a file manager.

There is much more to web content. You can create structures (pre-defined content formats), templates (designs for formatting structures), schedule when content is published and when it should be taken down (or reviewed), define related assets, stage multiple variations of your site, and more.

This is just the web content portion of the content management system. Liferay DXP is also great at managing file-based content.

1.2 Keeping Track of Documents, Images, Video, and More

Liferay's file-based content management system is in an application called *Documents and Media Library*. This application, as shown below, looks like the file manager that you're already familiar with from your operating system.

Like a file manager, you can browse files and folders in nested hierarchies. You can also mount other repositories that you might have in your environment, such as any system that implements Content Management Interoperability Services (CMIS). It generates previews of many document types. And, like a file manager, you can upload, copy, and move files between folders by dragging and dropping them. Of course, if you still want to use your operating system's file manager you can, because the Documents and Media library supports WebDAV, using the same credentials you use to log in.

Liferay DXP's Documents and Media library, however, is much more robust than a file manager is, because it's a full content management system. You can define ways of classifying files that may be of different types, but are meant for the same, overarching purpose.

For example, *metadata sets* are groups of fields describing attributes of a file. These can be combined to define *document types*. For example, you might create a document type called Meeting Minutes. The file format doesn't matter: whether it's a Microsoft Word document, an HTML file, or a text file, the document contains meeting minutes.

Of course, the system goes much further than this. Folders can be set so that only certain document types can be added to them. You can add workflow rules to folders so files go through an approval process that you define. In short, Liferay DXP's file-based content management system gives you what you need to manage and share files in a group.

There are many other ways you can use Liferay DXP, starting with its collaborative tools.

1.3 Using Liferay DXP as a collaborative platform

Liferay DXP includes a suite of collaborative applications, and they're all integrated together.

These applications range from personal productivity applications like a calendar and email, to community-building applications like message boards, blogs, and wikis.

This is a suite of integrated applications with all the features of similar, standalone applications. For example, Liferay DXP's Message Boards include categories and subcategories, message threads, captcha, RSS feeds, email notification, posting via email, and much more. But more than this, the applications are integrated with the rest of Liferay DXP's framework. Users log in and their profiles are used automatically by the message boards and all the other collaborative applications. And as you'll see later, functionality from the built in applications can be added to your own to provide features like comments in your own software, and you don't have to write any code to do it.

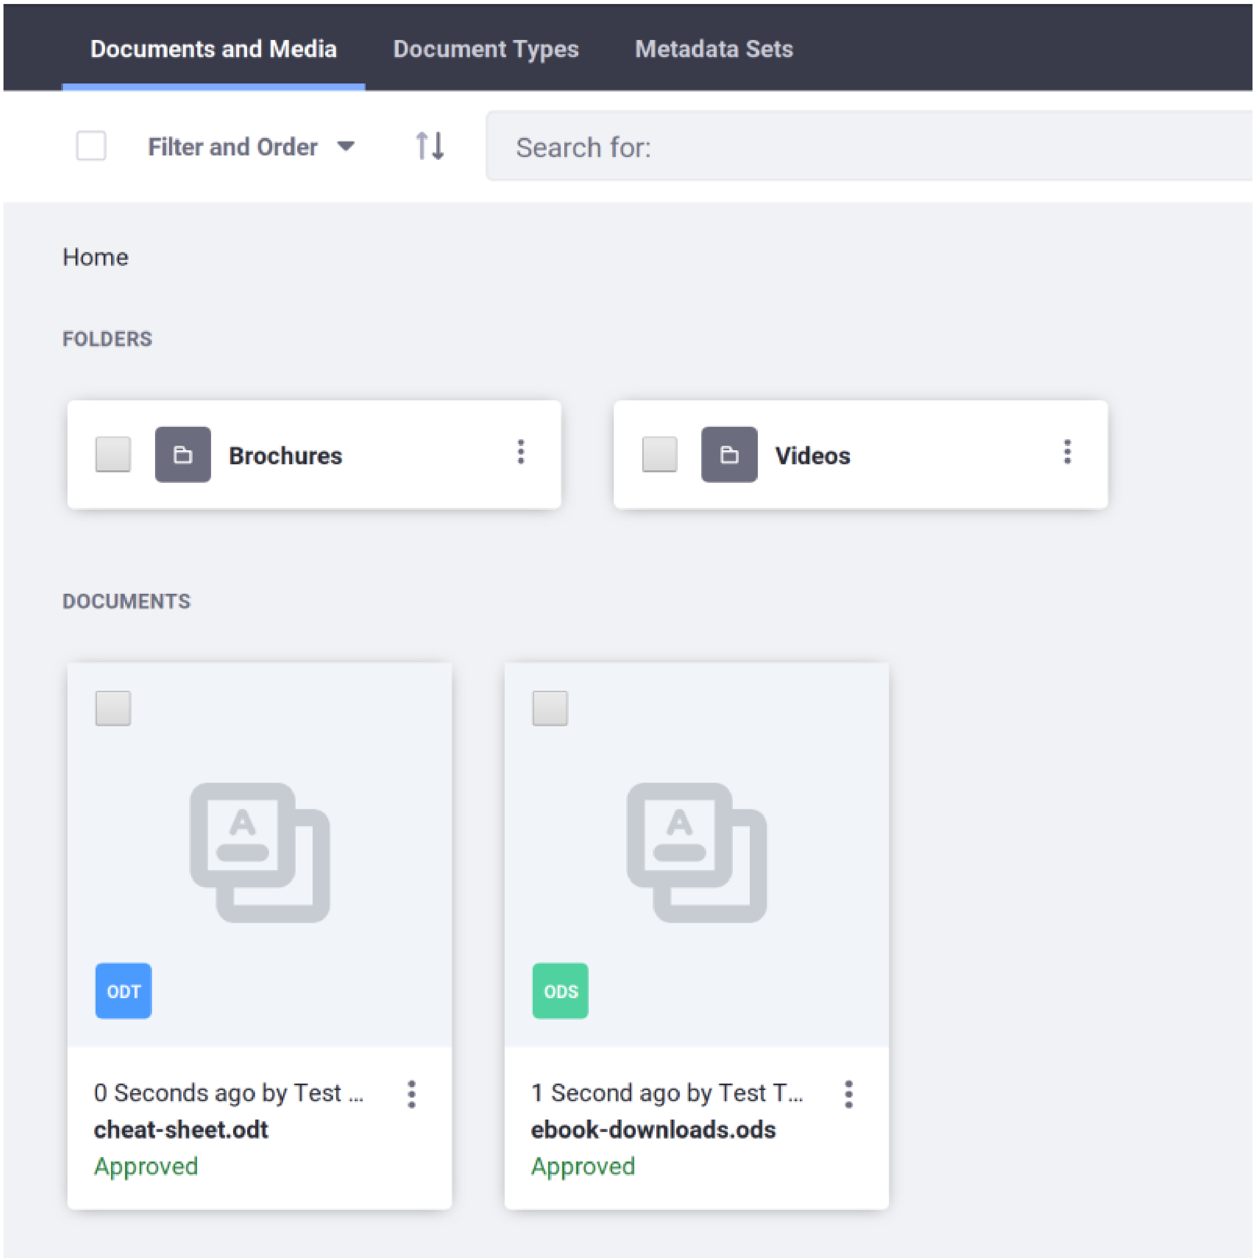

Figure 1.4: Liferay DXP's Documents and Media library was purposefully designed to be familiar to anyone who uses a computer.

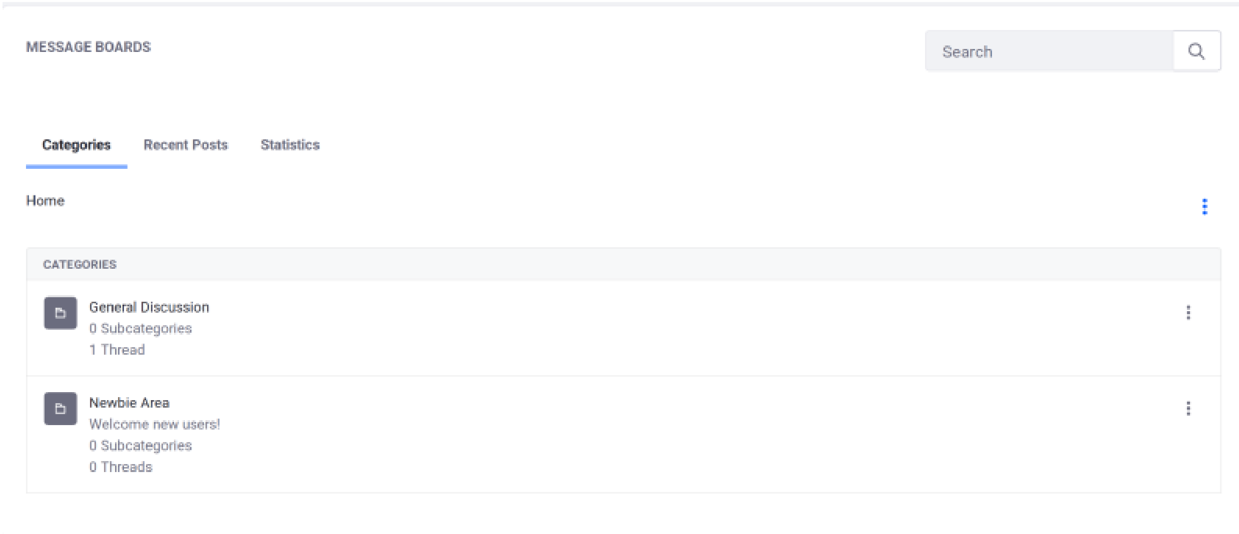

Figure 1.5: Liferay DXP's message boards are as fully featured as any standalone forum application, with the added benefit that they're integrated with the rest of the system.

One important feature of all the collaborative applications—as well as web content and documents—is the Recycle Bin. If users delete content that must be restored later, you don't have to find it in your backups: it's in the Recycle Bin.

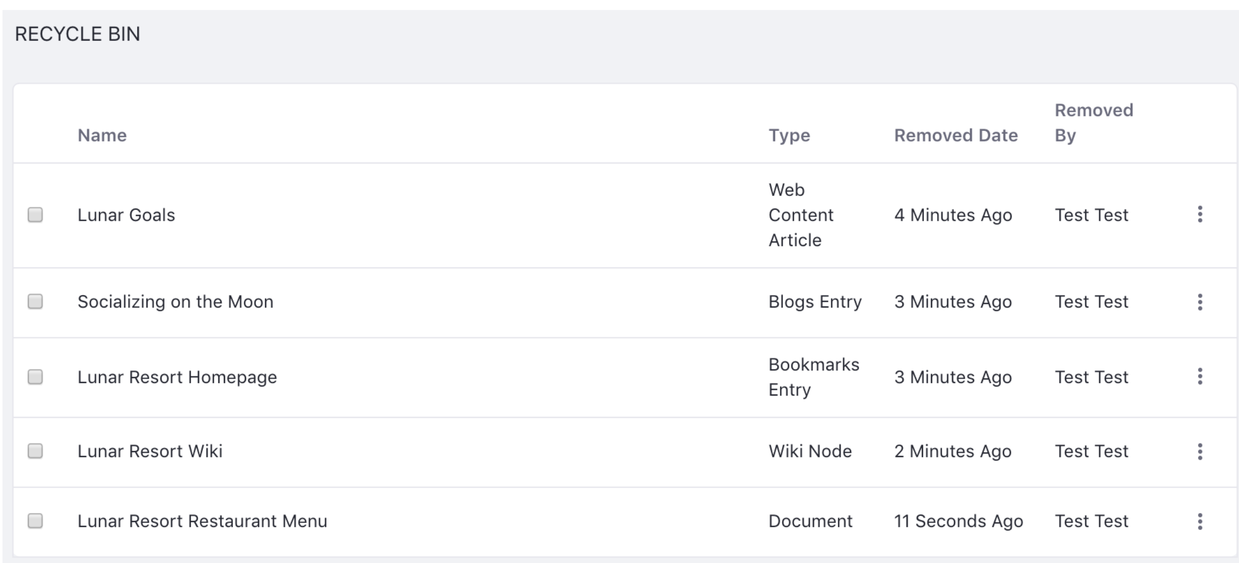

Figure 1.6: The Recycle Bin can hold any kind of content.

Liferay DXP's suite of collaborative applications includes a Blog (complete with blog aggregation features so you can publish multiple users' blog entries in one place), Message Boards, a Wiki, a Knowledge Base that you can use to publish a library of articles, activities display, and personal productivity applications like a calendar.

Liferay DXP includes what you need to enable users to collaborate. Next, you'll see how Liferay DXP's platform benefits developers.

1.4 Using Liferay as a Development Platform

Building your site on Liferay's development platform can give you a head start. It provides everything you need to support your applications, so you can concentrate solely on what *you're* building, and not the rest of the features your users expect to come along with it.

Imagine your application for a moment. Does it require users to register on your site? Can users comment on content contained in your application? Is there something that users can tag or categorize? If you think about how the code is organized, would it benefit from modularization? Do you plan on having multiple clients like web and mobile access a single back-end? Could you make use of a rich JavaScript framework with many components built into it? How about permissions—do you need to make information available to some users, but not to all users?

You get all this and much more. It's a very powerful platform, and certainly worth your investigation. If you're a developer, it behooves you to check out our developer documentation.

A Great Integration Platform

If you're building an enterprise system, portals were designed in the first place to be a single point of entry to your users' applications and content. Since Liferay DXP integrates with user directories such as LDAP and Active Directory, single sign-on systems such as SAML and OpenSSO, and authorization systems like OAuth 2.0, it fits well into your enterprise systems. This makes it a great integration platform for existing applications.

Liferay DXP, since it adheres to the JSR standard for portlets, was designed from the ground up for application integration. You can add any application installed on the system to any page. You can make use of APIs provided by other systems to integrate their data into an application window in Liferay. And applications you create with Liferay's Service Builder framework can be web service-enabled from the start.

Developing applications is only one part of the platform. Many users value Liferay DXP for how it can be customized. Next, you'll see how you can customize Liferay DXP so that it looks and operates exactly the way you've envisioned for your site.

1.5 Extending and customizing Liferay for your own needs

Liferay DXP was specifically designed to be extended and modified, beyond what even most open source projects provide. Though its source code is available, Liferay DXP's developers have created many extension points so you don't have to customize the source to make Liferay DXP yours.

The first (and easiest) way of customizing parts of Liferay DXP is with Application Display templates. These let you change the way built-in applications look. For example, if you don't like the Documents and Media Library's file manager view with large icons, you can create an Application Display template that shows documents in some other view. If you don't like the layout of the Blogs application, you can change it so that it has the look you want.

Liferay DXP goes far beyond this, though. Special software components called *modules* enable developers to change Liferay's interface and behavior without having to modify any of Liferay DXP's source code. This provides you all the benefits of building your site from scratch, but without all the effort to actually build from scratch. If you want to make a change to the user registration screens, add support for a proprietary single sign-on mechanism that you've written, add a feature to the message boards application or anything else, you can make those customizations. And if you're a developer, you know that it's a whole lot easier to customize something that almost does things the way you want than it is to write that feature from scratch.

1.6 Conclusion: What Liferay DXP Really Is

So what is Liferay DXP? As you can see, it's hard to describe, because it does so much. What we've essentially done is say it's a totally awesome content and document managing, user collaborating, socially enabling, application developing, corporate integrating, completely customizable platform for building the Internet. If we'd said that up front, you may have thought we were exaggerating. Hopefully now you can see that we really believe it, and we continue striving to make Liferay DXP better and better.

But please don't take our word for it. Download the open source, Community Edition or a trial of the supported Digital Experience Platform yourself. See how it can solve some of the problems you may be facing. If you have questions, all the documentation is right here, you can ask our friendly community for help, or you can contact a sales representative.

Chapter 2

What's New in 7.0

7.0 launches new features for efficient, beautiful page design, along with an extensive upgrade to forms. The user experience has also been streamlined throughout the platform to help save time while managing sites. There are new tools to support data protection in accordance with the recent GDPR requirements. Finally, there are several significant technology updates that help you develop new competitive advantages with modern development tools for mobile and web.

Evolve your digital presence with new features to create stunning, personalized experiences for every audience.

2.1 Page Creation

7.0 introduces a powerful new way to design websites. From carefully designed page fragments to full control over page menus, Liferay DXP frees web developers and designers to execute web experiences exactly as they envision. This is an enhancement of the content management system that you can adopt as you see fit. To differentiate the two, the new pages are called Content Pages, while the existing system uses Widget Pages.

Content Pages

Users can now easily create and add unstructured content directly to pages. This is useful for site pages that don't need to leverage structured web content, such as one-off landing pages for marketing campaigns.

When creating a new page, users can choose either the new content pages or widget pages, which use the traditional method of creating pages through adding and configuring applications.

Fragments

Fragments are a new way of creating and implementing content page designs. Web developers can now save page sections as Fragments and reuse them across a Site. Developers can now create a library of designed components for others to build pages quickly without having to touch code. Fragments are organized in collections and leverage familiar asset management features such as drafts, thumbnail previews, search, and permissions.

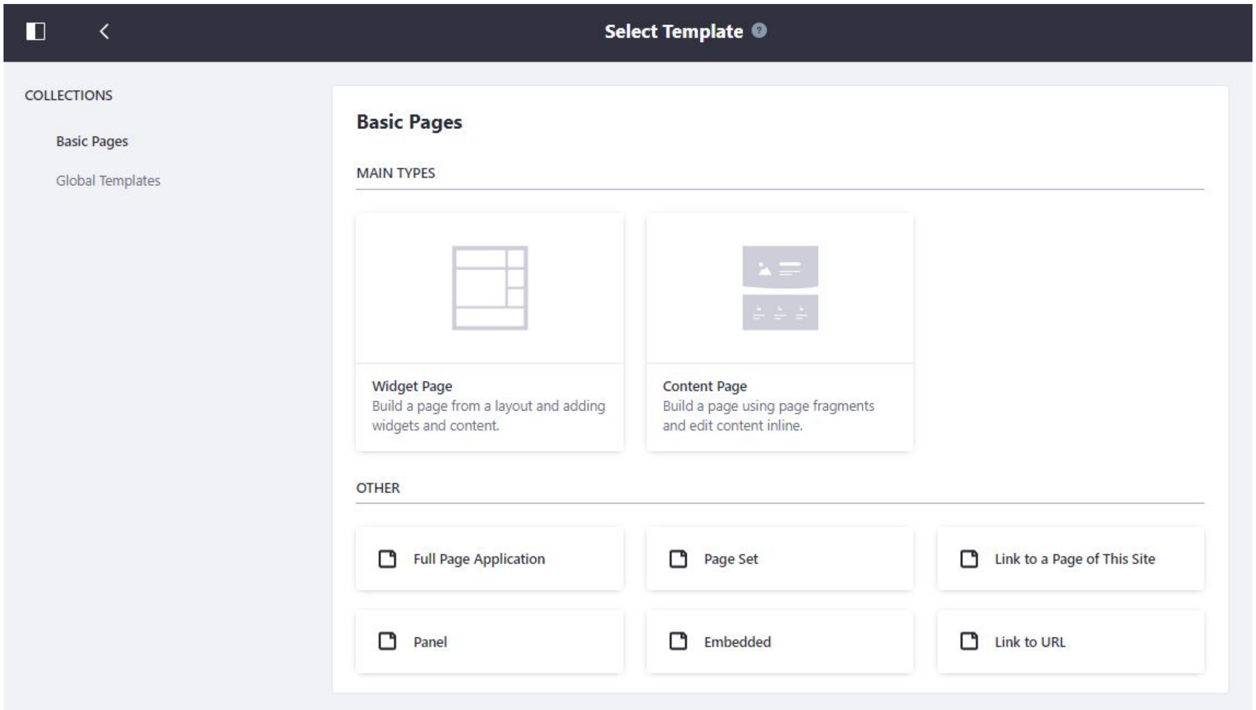

Figure 2.1: Users can now choose to create Content Pages.

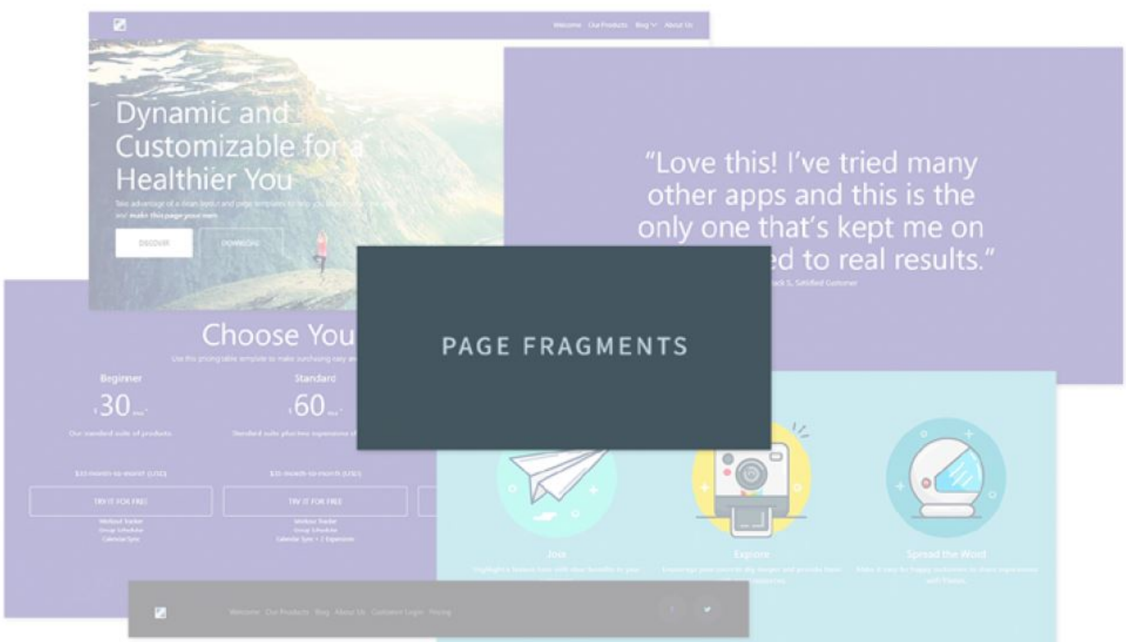

Figure 2.2: Non-technical users can use Fragments as building blocks for pages.

Fragment Editor

Web developers can use the Fragment Editor in the browser to create or edit their Fragments. Alternatively, developers can create Fragments with their preferred tools and import the fragments into Liferay.

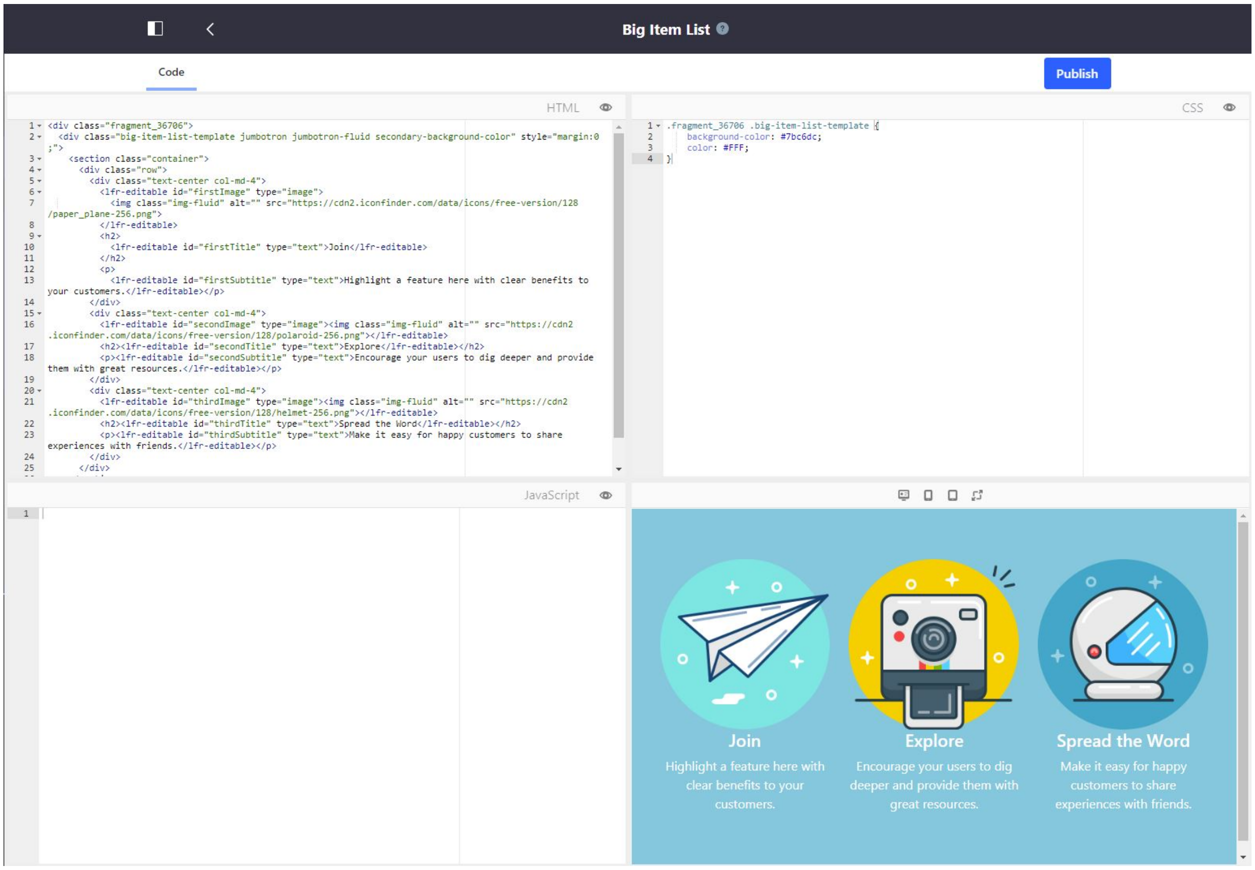

Figure 2.3: The Fragment editor has a three-pane view for HTML, CSS, and JavaScript, with a preview pane rendering the final result.

Page Editor

The introduction of Fragments opens up new ways to create pages and templates through a visual page editor.

The new Page Editor is for laying out page designs visually and saving them as reusable templates. Users can search through collections and easily add, remove and position fragments on the page. Marketers can then customize text with in-line editing, swapping in new images and other elements. Fragments can display Liferay's out-of-the-box applications, and you can configure them within the page editor.

Display Pages

Display Pages have been improved to make it easier to create standard templates for web content that must have a consistent look and feel, such as press releases. Use page Fragments to implement

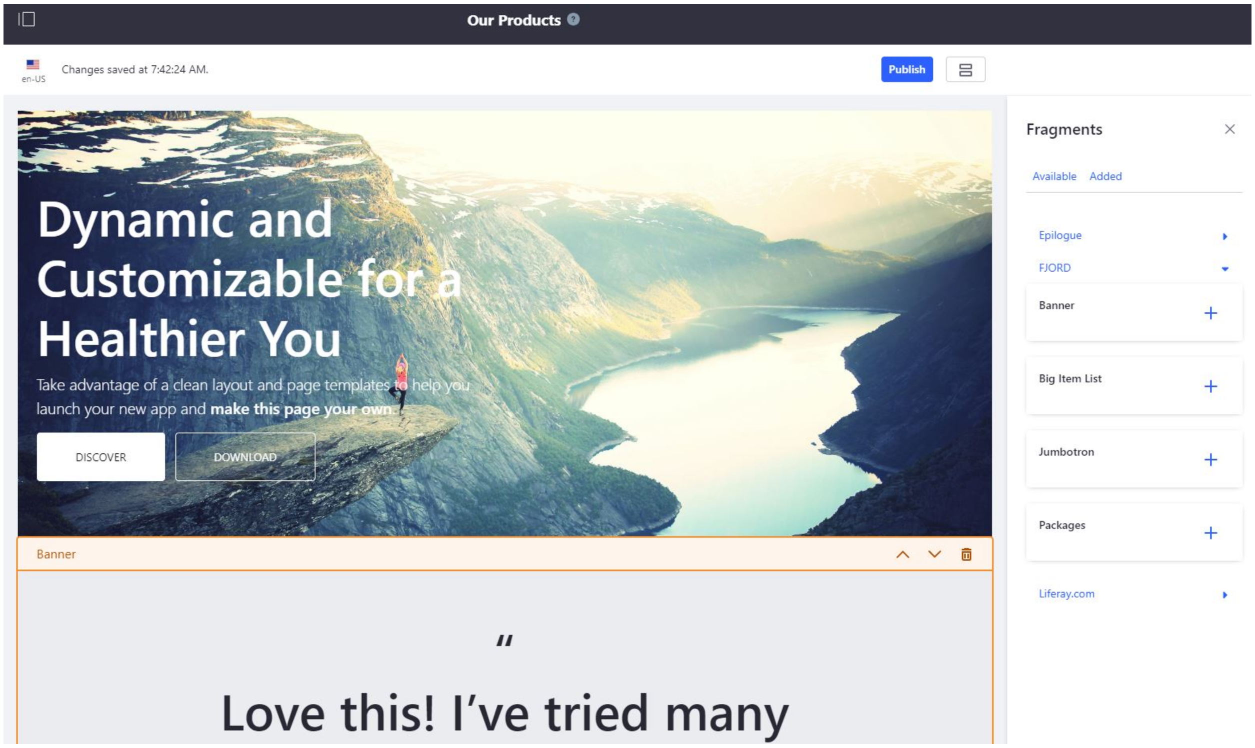

Figure 2.4: The Page Editor makes it easy to build pages out of Fragments.

designs and map content sections without touching code. When web content is published with a display page template, it automatically gets its own page with a unique URL, replacing the default content in the template with the newly added web content.

Menus

Menus are now decoupled from page navigation. Now you have the freedom to create custom menus for sections of the site or remove marketing landing pages from menus. You can easily manage menu hierarchies and save different menu sets with a new drag-and-drop interface.

Forms

Forms have extensive new functionality, including a set of conditional rules that make forms dynamic: if something happens in a field, an action elsewhere on the form can be triggered. This lets you create calculations, offer follow-up questions based on responses, and much more. Forms are now localized: they can be translated into any language. There are new fields and properties, form fragments, auto-saving of forms, and so much more it had to be described separately.

User Experience

7.0 rolls out several user experience refinements for externally facing content such as blogs, as well as omni-channel support for how media is displayed and delivered on different devices. These integrated improvements help you deliver a better experience to your users out of the box.

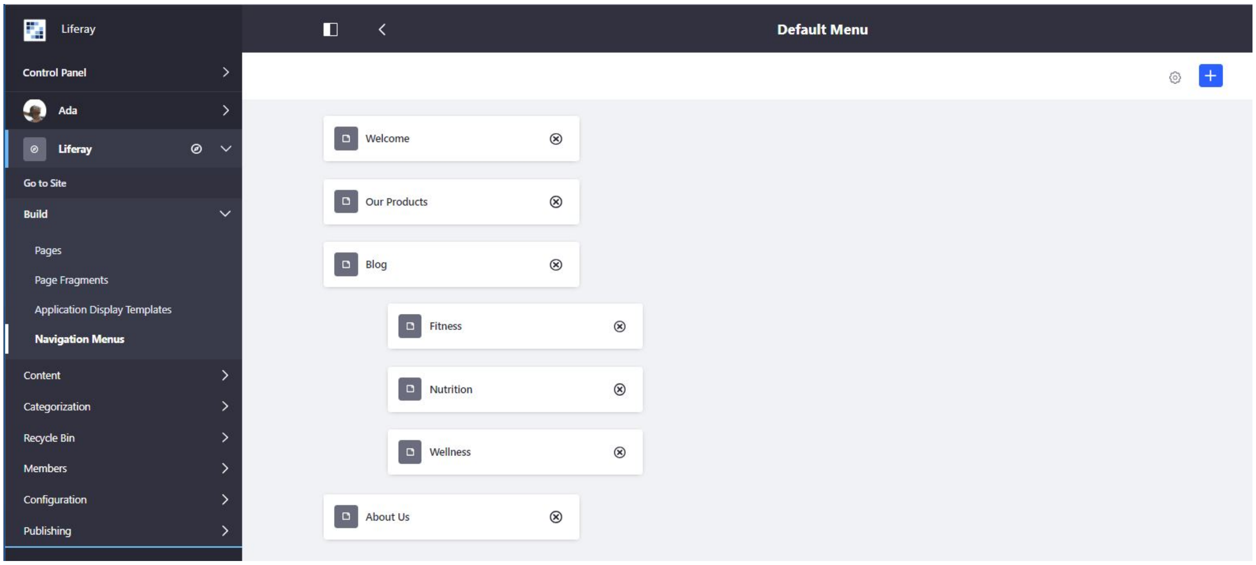

Figure 2.5: Menus are now decoupled from page navigation, and can be edited independently.

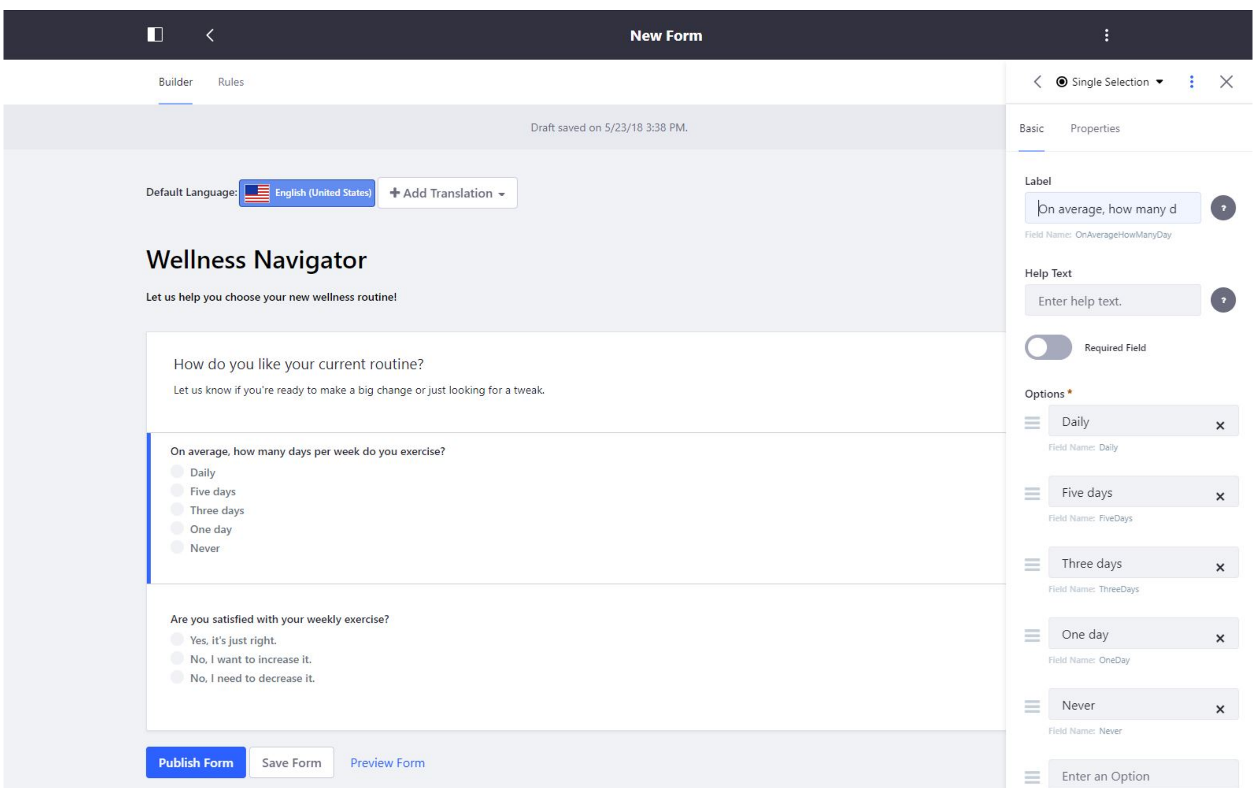

Figure 2.6: Forms have many improvements.

Adaptive Media

Adaptive Media dynamically adjusts images to best fit the screen size of the device being used. It also offers deep control over how images are loaded and displayed, which helps you address performance issues across a wide variety of devices and varying network speeds between users and countries. Adaptive Media does most of the work in the background automatically, but developers can edit image resolutions and define devices that trigger various resolutions. This level of control enables consistent experiences that avoid poor page layouts and slow load times.

Blogs and Message Boards

Blog refinements make it easier to deliver experiences that are tailored to your blog audience. These include support for creating friendly URLs, displaying estimated reading times, unsubscribing from email notifications, a new cards design, and support for videos from external services. Message Boards now supports drag and drop for uploading attachments, section renaming, category and thread grouping, notification management, and a new design for comments.

2.2 Administration Improvements

In addition to its new features, 7.0 is easier to administer. Streamlined administration tools mean less time managing sites, while still offering granular control.

OAuth 2.0

OAuth 2.0 has become a de-facto standard that allows users to authorize access to parts of their accounts without giving up authentication credentials. With 7.0, now users can "Sign in with [insert your site here]," granting mobile and web applications secure, token-based access to user profile information they have complete control over and can revoke at any time.

Data Protection

7.0 introduces new data protection tools to help companies address GDPR regulations and maintain control over how their platform manages user data.

You can now erase a user's personal data and export a user's personal data in a machine-readable format upon request. For data erasure, administrators can review content that potentially contains personal information and edit or delete as needed through a simple interface. Both tools include APIs for third-party apps to implement this feature or override the default behavior for out-of-the-box apps.

Search

7.0 improves search administration and uses Elasticsearch 6 as the default search engine, giving users more options for implementing and managing enterprise search for their sites.

A new Control Panel makes it easier to take care of all administration tasks with the click of a button. Users can configure the search engine, start and monitor re-indexes, and much more.

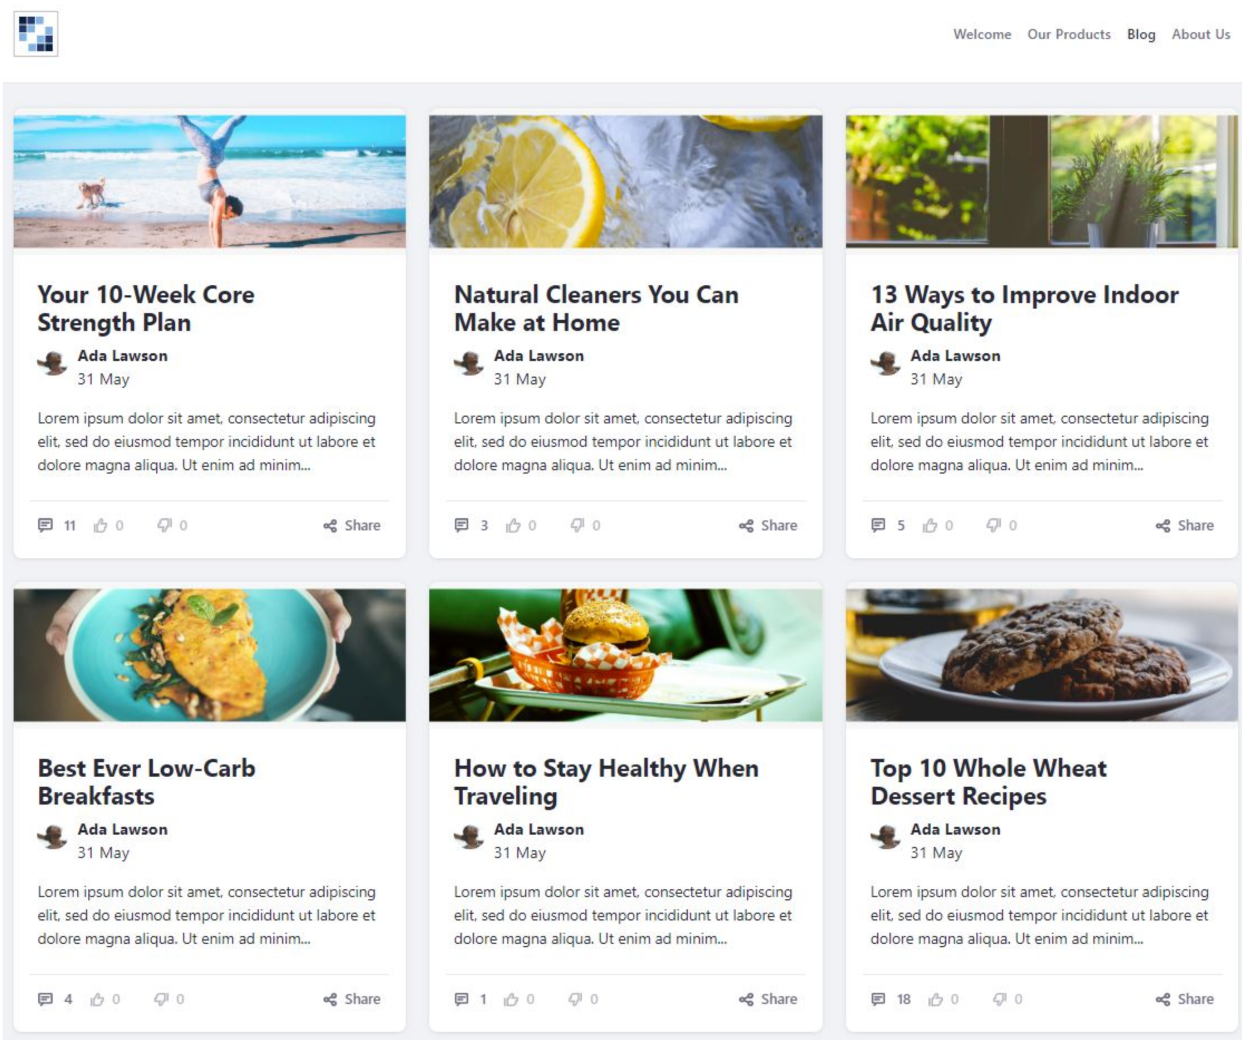

Figure 2.7: The new cards design for blogs displays entries in a visual grid.

Page Management

A new interface for visualizing and managing complex page hierarchies makes page management easier. Page templates and display pages are now integrated nicely, bringing page management into one central location.

Workflow Management

Workflow management has received a complete UI overhaul, with all configuration consolidated under one area in the Control Panel. Existing workflows can now be duplicated, definitions are versioned, and you can save drafts and restore previous versions.

2.3 Developer Improvements

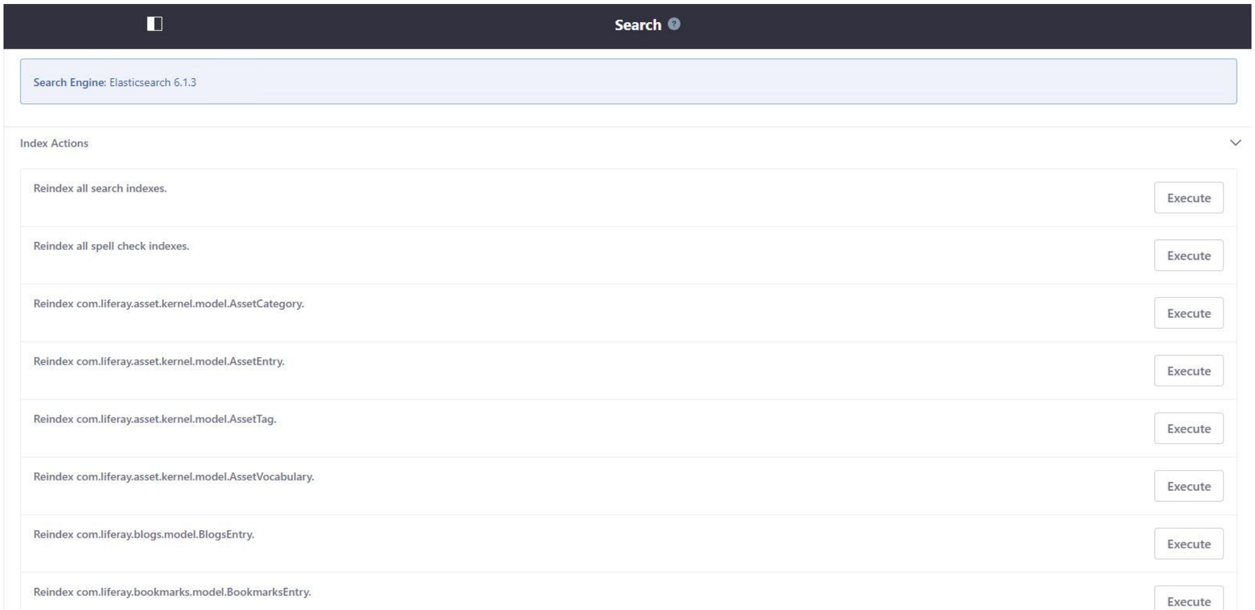

Figure 2.8: Search administration is now separated from server administration.

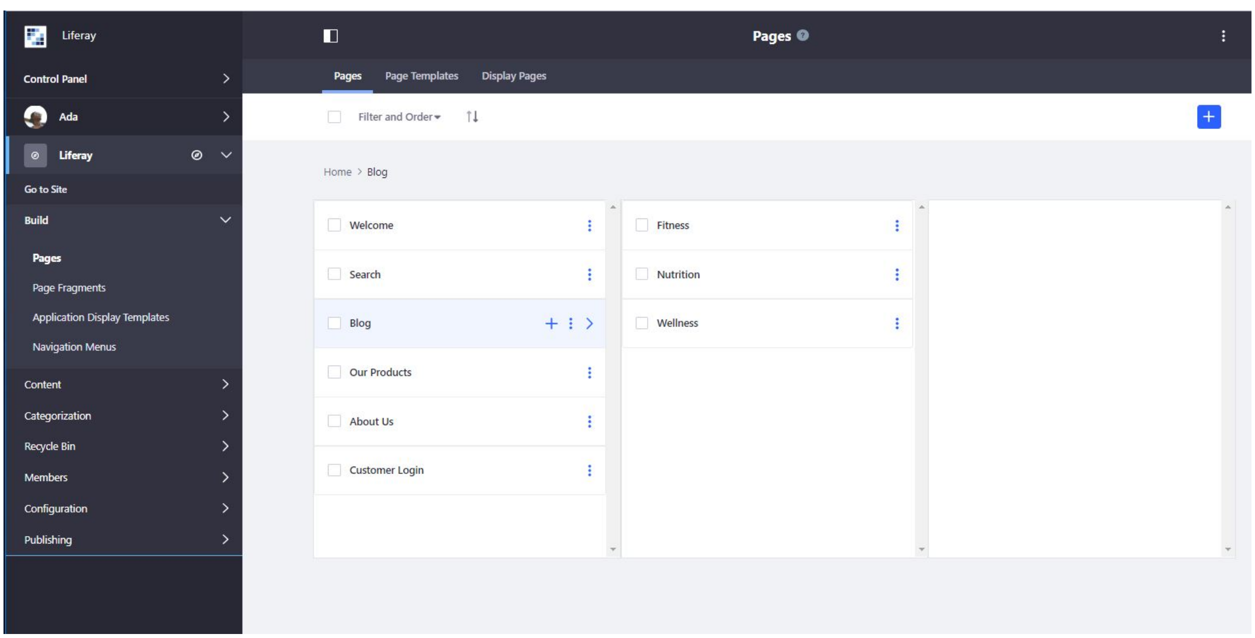

Figure 2.9: The new page management interface puts all page functions in one place.

7.0 includes an updated collection of tools to facilitate the support and development of Liferay projects.

Targeting a Liferay Platform

Liferay Workspace helps target a specific release of Liferay DXP, so dependencies get resolved properly. This makes upgrading your applications easy: specify your target platform, and Workspace points to the new version. All your dependencies are updated to the latest ones provided in the targeted release.

Resolving Modules Before Deployment

Avoid the painful process of deploying modules only to be met with console errors or mysterious problems by resolving modules before deployment. This can be done by calling the new *resolve* Gradle task provided by Liferay Workspace.

7.1 Upgrade Planner

The Upgrade Planner in Liferay Developer Studio helps you upgrade your legacy application code to Liferay DXP:

- Identifies code affected by the API changes
- Describes each API change related to the code
- Suggests how to adapt the code
- Provides options, in some cases, to adapt code automatically.

IntelliJ Support

Liferay development is now officially supported on IntelliJ IDEA which offers wizards to

- Create a Liferay Workspace
- Create projects leveraging Liferay's project templates
- Create a Liferay server runtime for project deployment and debugging

Maven Support for Blade CLI

Create Maven projects and Maven Liferay Workspaces using Blade CLI.

Hybrid Mobile App Development

Liferay Screens 3.0 enables software developers to use Apache Cordova or Xamarin to build cross-platform applications from one codebase designed for the web and embed that content into a Screens app for mobile use. Sites and applications designed for PC can be rendered in screenlets with no additional code. The resulting apps allow native mobile capabilities and navigation to be mixed with HTML content seamlessly.

Modern JavaScript Frameworks Compatibility

Liferay DXP leverages its own npm bundler so developers can manage dependencies between applications. 7.0 provides support for popular JavaScript frameworks such as Angular, Vue.js, React and modern JavaScript workflows, so that npm modules can be deployed inside of Liferay DXP.

Modularity Update

New search applications such as Search Results, Search Bar and Category Facets allow for greater flexibility in page construction. These applications come from the previous Search application, but have been divided into more useful components.

The message boards services have also been modularized and extracted out of the core, making it easier to manage and update Message Boards independently.

Chapter 3

Web Experience Management

Experience: consider that word for a moment. Not the type of experience you gain with repetition, but the contact or encounter you have with something. Suppose you're buying a new phone. It's easy to look at the phone with the biggest screen or the fastest processor and say that it's the best, but what will your experience be? The "best" phone might not be the fastest or the biggest; it might be the phone with the longest battery life or the most comprehensive suite of integrated apps. Or it might not be any of those things. Experience isn't always something that you can quantify with specs or features.

Liferay takes the experience factor very seriously when it comes to site and content management. The Web Experience Management suite is focused on providing the best experience for users building websites. When it comes to web experience, just like with your phone, everyone is looking for something different. Some smartphone users might love watching videos during their commute with a big beautiful screen, while on the other side, some people might be more excited about a small sleek phone with a great battery life that fits easily in their pocket and simply does all the basic communication they need.

Liferay's Content Management is the big beautiful phone and the sleek utilitarian one all in one. Marketers will find easy to use tools to build content without having to write any code or peak under the hood. Developers will find powerful tools like Structures and Templates that enable them to create dynamic content. And designers will love how Fragments and Content pages provide a way to perfectly realize their designs.

3.1 Authoring Content: Structured and Inline Content

The primary goal of Content Management isn't to show off the flashiest new features or follow all the latest trends in design, but to provide you with the tools you need to create digital content that communicates your message clearly and effectively. With this in mind, Liferay offers two core approaches to help you accelerate and simplify creating and organizing content: Structured Web Content and Inline Content.

Structured Web Content

If you've entered content into a CMS before, you may be familiar with the process of filling content into various fields like this:

- Title
- Abstract
- Text Body

This is an example of Structured Content. Structured Content is created within a predefined content structure and then added to pages as needed. The structure defines the fields and then a template defines its styles. The content is then saved, ready to be added to a page later.

The structure defines what kind of content you are creating and provides different types of fields that can be used. A developer could create a format for publishing articles that contains a **Title**, **Header Image**, **Body Text**, and a **Key Quotation**. The template defines how the elements of the structure are rendered. You could style the elements in the structure with the **Title** as large bold text, the **Header Image** as a full page width block above the title, the **Text** as standard text, and the **Key Quotation** as large font italics with a thin border that displays within the main text section. A content writer or marketer could then create any number of articles, all having a uniform style based on the structure.

In Liferay, those articles could be added to pages across the site, or displayed dynamically with tools like the Asset Publisher and Web Content Display Pages.

Inline Content

Inline Content is content that is created directly within a page. Rather than filling in fields to create content that's added to a page later, you have a completed page design where you edit the text and image content. Content Pages start with a design which is then created with Fragments. Inside the Fragments, a developer can define where text and images can be placed or edited. Marketers and content writers are then free to write or add images within the page and publish it.

With content pages, basic HTML and CSS define the primary design, while JavaScript and Liferay specific tags can add dynamic behavior to the Content Page. After a developer creates the page and a content writer or marketer provides the content, the content exists inline within the page, and is published with it.

What's best for your use case?

When you step back to look at the big picture, what you see are two different paradigms for building pages: content pages, where the content is built into the page; and widget pages, where content and other features can be added, removed, and rearranged as desired.

Often it is helpful to have reusable elements or content that can be moved around a page or placed anywhere on a site. For example, you might have a content based banner which you want to be able to drop onto any various pages with different layouts and styles, or you might have content that uses the same template to create similar items for different pages. Structured Content on widget pages is the tool you need to quickly create what you need and manage these cases. Widget pages with structured content are great for some cases, for example:

- "Portal" pages, where you provide users a gateway into several different services or providing aggregated information.

- Pages where the primary focus is widgets.
- Pages that are based around structured content.

In other cases, you need to create a page as a complete unit. For example, you have a series of marketing driven landing pages that must match a specific design and have associated content intended for use on that page or with that campaign. Content Pages provide the best tool for quickly bringing a design to life and empowering marketing with inline content. Content Pages are useful for pages like:

- Landing pages
- Front pages that provide marketing information or a direct path into the website.
- Pages with multiple variations and small graphical or textual changes across a large number of pages.

Most sites need a little bit of both. Read on to learn more about building sites with Liferay and how Content Pages and Structured Content can help you do that.

Chapter 4

Building a Site

A site is a set of pages where content or applications are published. Sites can be independent or serve as an associated organization's website. You can create as many different sites as you like within the context of a single Liferay instance.

You'll start with a tour of the site management user interface. Then you'll create a custom Lunar Resort Example instance and explore ways to create sites and pages for that Liferay instance. Finally, you'll learn how to change various settings for sites and pages to meet your needs. To begin building a site, continue on to the next section.

Chapter 5

Site Management

You can have many Sites on one Liferay instance, which work together to create one complete website, or you can simply have one Site which contains all of your pages and content—or anything in between. In this section you'll look at the interface for creating and managing Sites, create a Site, and learn how to use Site Templates for more efficient Site creation.

5.1 Understanding Site Management

Whether you're building a large corporate webSite or a small Site for facilitating collaboration among team members, supporting different kinds of collaboration and social scenarios is a must. Liferay's Sites provide three membership types:

Open: Users can become members of the Site at any time.

Restricted: Users can request Site membership but Site administrators must approve requests for users to become members.

Private: Users cannot join the Site or request Site membership. Site administrators must manually select users and assign them as Site members.

In addition to these memberships, when a Site is associated with an organization, all the users of that organization are automatically considered members of the Site.

You can view all the available open and restricted Sites by adding the My Sites application to a page and accessing the *Available Sites* tab. You can request access to any of the Sites you're not already a member of by selecting the Site's *Options* button

( ) and clicking *Join*.

Site Scope

Members of a Site can be given additional privileges in the Site by using permissions. It is also possible to assign different roles within the Site to different members. This can be done through *Site Roles*, which are defined equally for all Sites or *Teams* which are unique for each Site. These concepts are discussed later.

Liferay DXP separates Site-scoped information from the Control Panel by placing it in the Site menu. From this menu, you can select the specific Site to work on. The Site Administration panel is available for your Site, which includes Build, Content, Categorization, Recycle Bin, Members, Configuration, and Publishing.

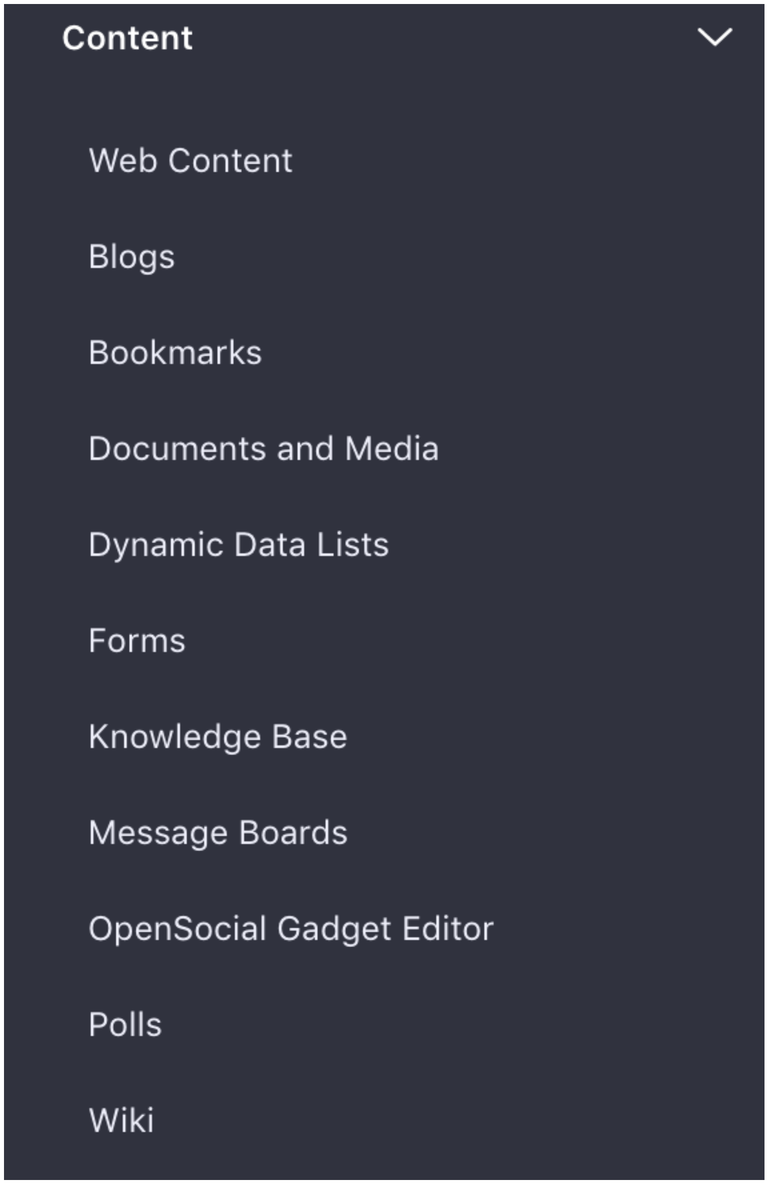

Figure 5.1: Your Site's content resides in the Site Administration menu.

Site Hierarchies

Sites can also be organized hierarchically, just like Organizations. The difference between Sites and Organizations, of course, is that Sites organize pages, content, application data, and users (via Site memberships) whereas organizations only group users. Content sharing is available for Sites within the same hierarchy. For instance, if a parent Site has a document type called *Lunar Presentation* and all its child Sites should have a copy, the parent Site's administrator can enable content sharing to share the document type automatically with its child Sites. Also, content sharing privileges can be set to let every Site administrator share content across Sites they manage. You can share the following content across Sites:

- Web Content Structures
- Web Content Templates
- Document Types
- Vocabularies and Categories
- Widget Templates
- Data Definitions (Dynamic Data Lists)

Please refer to the Sites Admin Portlet section of Liferay's `portal.properties` file for a list of relevant configurable properties. For example, the `Sites.content.sharing.with.children.enabled` property can disable content sharing between Sites and child Sites, disable it by default while allowing Site administrators to enable it per Site, or to enable it by default while allowing administrators to disable it per Site.

The Sites Directory application is a configurable app that shows a hierarchy of Sites and child Sites. It enables users to navigate to any of the displayed Sites. To use this app to display Site hierarchies, add it to a page, open its Configuration window, and under Display Style, select *List Hierarchy*. The My Sites Directory application is similar to the Sites Directory application, except that it lists only the Sites a user belongs to.

Each child Site in the hierarchy has its own administrator, and the Site Administrator role permissions do not flow down to child Sites in the hierarchy. If a Site Administrator creates a child Site, he or she has the same permissions in that child Site. This is not, however, because of inheritance. It is only because creating a Site makes you the Owner of that Site. A Site Administrator or a parent Site has no default role in any child Sites created by other Site Administrators.

If you wanted a user to have administrative access to all Sites in a Site/child Site hierarchy, you must create a role based on the Site Administrator role that has the permission *Manage SubSites*.

The Site Map application helps users navigate a Site. A Site administrator can configure a root page and a display depth. Just as Sites can have hierarchies, so can the pages within a Site. The display depth of the Site Map application determines how many levels of nested pages to display.

Site Members

Another useful administrative application is the Site Members application. This enables administrators to survey all the users, organizations, and user groups that reside in the Site. Similarly, Liferay provides the Portal Directory application, which functions the same as the Site Members app, but globally scoped for all Sites in the instance.

Page Sets

Sites have two categories of pages called page sets. There are two kinds of page sets: public pages and private pages. A Site can have only public pages, only private pages, or both. Private pages can only be accessed by Site members. Public pages can be accessed by anyone, including users who haven't logged in. It's possible to restrict access to pages at the page set level or at the level of individual pages through the permissions system. Public pages and private pages have different URLs and can have different content, applications, themes, and layouts.

Building a corporate intranet is a typical use case for Sites. A corporate intranet could have Sites for all the organizations in the company: Sales, Marketing, Information Technology, Human Resources and so on. But what about the corporate health and fitness center? That's something

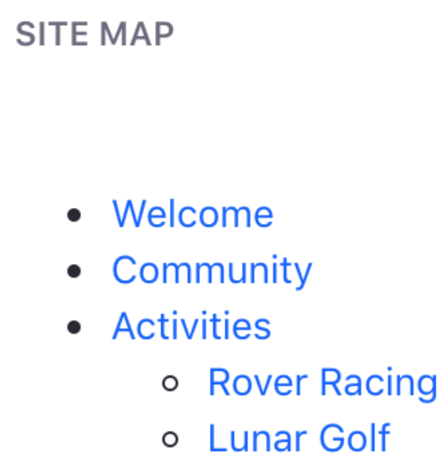

Figure 5.2: The Site Map application lets users navigate among pages of a Site organized hierarchically.

everybody in the company, regardless of organization, may want to join. This makes it a good candidate for an open and independent Site. Similarly, the home page for a corporate intranet should probably be placed in an open independent Site so any member of the instance can access it.

For other kinds of websites, you may want to use independent Sites to bring users together who share a common interest. If you were building a photo sharing website, you might have independent Sites based on the types of photos people want to share. For example, those who enjoy taking pictures of landscapes could join a Landscapes Site and those who enjoy taking pictures of sunsets could join a Sunsets Site.

There is always one default Site, which is also known as the main Site of the instance. This Site does not have its own name but rather takes the name of the instance. By default the instance name is *Liferay* but this value can be changed through the configuration of the setup wizard. The instance name can also be changed at any time through the Control Panel within *Configuration* → Instance Settings*.

5.2 Adding Sites

Sites can be created through the Control Panel by a Liferay administrator. The Control Panel provides an administrative interface for managing your Liferay instance. There are four main sections of the Liferay Control Panel: Users, Sites, Apps, and Configuration. In this section, you'll learn how to use the Control Panel to manage Sites. For information about the Apps, Users, and Configuration sections of the Control Panel, see the Using the Liferay Marketplace, Managing Users, and System Wide Settings sections, respectively.

Tip: If you're signed in as an administrator, you can access all Sites by navigating to the Site Administration menu from the Control Panel. To manage a single Site, navigate to the Site by going to the Menu and clicking the *Site Selector* button () from the Sites dropdown menu and selecting the appropriate Site name. Once finished, the Site administration options (i.e., Navigation, Content, Members, etc.) for that Site are available.

Now, you'll add a Site for the Lunar Resort.

1. Navigate to the Control Panel and select *Sites* → *Sites*.
2. Click the Add icon () at the top right of the page.
3. Select a *Blank Site*.

   Any available Site templates appear for you to select. Site templates provide a preconfigured set of pages, applications, and content that can be used as the basis of a Site's public or private page set. To create a Site from scratch, select *Blank Site*. Otherwise, select the name of the Site template you want to use. If you opt to create a Site from a Site template, you have to choose whether to copy the Site template's pages as your new Site's public or private page set. If other Site templates are created, they will appear in the Add menu as they become available.
4. Name your Site "The Lunar Resort"

After you enter the name, you will be prompted to enter additional information about the Site and configure certain Site settings.

Name: names the Site you wish to create. You also have the option to translate the name for many different languages. This can be done by selecting the language flag under the Name field, and inserting the name in the selected language. Liferay saves the name translation for each language and displays the translated Site name when that specific language is selected for the instance. If a name translation is not provided, the default instance language's name is displayed.

Description: describes the Site's intended function. The description can also be translated to other languages; see the Name description for more information on translating the Site's description.

Active: determines whether a Site is active or inactive. Inactive Sites are inaccessible but can be activated whenever a Site administrator wishes.

Membership Type: can be open, restricted, or private. An open Site appears in the My Sites app and users can join and leave the Site whenever they want. A restricted Site is the same except users must request membership. A Site administrator must then explicitly grant or deny users' requests to join. A private Site does not appear in the My Sites app and users must be added to it manually by a Site administrator.

Allow Manual Membership Management: determines whether to allow or disallow users to be manually added or removed from the Site. By default, manual Site membership management is enabled. This allows administrators to manually assign users to the Site. It also allows users to join open Sites or request membership from restricted Sites using the My Sites app. For organization Sites, manual Site membership management is disabled, by default. This causes organization members to be automatically assigned membership following the organization's membership policy. Also, because manual membership management is disabled for organization Sites, by default, the *Users* section of *Sites* is unavailable. To activate the *Users* functionality for your organization Site, you'll need to check *Allow Manual Membership Management* after creating the organization Site by navigating to its *Site Settings* menu.

Note: It's possible for Site memberships to be handled automatically by a membership policy. The membership policy can check various pieces of information from each user, such as their first names, last names, birthdays, job titles, organizations, and user groups. Using this information, the Site membership policy can automatically assign members to the Site. If your Site will implement a membership policy, your Site administrators can disallow manual membership management for their Site. When the Allow Manual Membership Management option is disabled, the *Members* section of Site Administration (Site Memberships and Site Teams) is hidden, even from administrators.

Parent Site: lets you select a parent Site for the Site that's being created. Sites can be organized hierarchically. Using hierarchical Sites provides a simplified way to manage Site memberships and Site content sharing. For organizations that have attached Sites, the organization hierarchy should match the Site hierarchy. When you select a parent Site, an additional option appears: *Limit membership to members of the parent Site*. If this option is enabled, the Site's membership policy performs a check so that you can only assign members to the current Site if they're already members of the parent Site.

2. Set the *Membership Type* as *Restricted*.

3. Leave the remain defaults and click *Save*.

When creating a blank Site or organization Site, the Site is not immediately viewable. This is because Sites without a page are impossible to view. Therefore, before you can view your Site, you must first create a page for it. To add a page for your temporarily invisible Site, navigate to the *Navigation* option from Site Administration. Then add a public page. After adding your Site's first page, it renders and your Site is viewable. For more information about adding pages, see the Creating and Managing Pages section.

You can also categorize your Site template using tags and categories by selecting the *Categorization* menu from the bottom of the page. To learn more about using tags and categories in Liferay, see the Organizing Content with Tags and Categories section. Lastly, at the top of the page is an additional tab named *Social*. This tab manages whether users of your Site can mention other users. You'll learn about mentioning users later in the Social Collaboration sections.

When creating a Site from a Site template, you're asked if you want to copy the pages from the template as public pages or as private pages. By default, the Site is linked to the Site template and changes to the Site template propagate to any Site based on it. A checkbox appears for unlinking the Site template if the User has permission to do so.

Once the Site has been created, you should configure its settings to fit your needs. You can learn more about Site Settings in Configuring Sites.

Next, you'll learn about creating pages.

Chapter 6

Adding Pages to Sites

In the previous section, you learned how to create sites. You may have gathered from that section that sites aren't particularly useful without pages. In fact, sites primarily exist for the sake of organizing pages and content, so now you'll learn about the different types of pages in Liferay, and how to select the best tools based on your use cases. You'll also learn how to manage pages and use various configuration options.

6.1 Creating Pages

After you create a Site, you can add new pages and maintain them. You can do everything you need with pages from Site Administration.

1. If you're not currently on the Site you want to edit, click the *Site Selector* button () next to your current Site name in the Menu and select your desired Site.
2. Go to *Site Administration* → *Build*.
3. Click on *Site Pages*

From here, you'll create pages and page templates.

Note: Pages are always part of page sets, and page sets are always associated with Sites. Even users' personal pages are part of their personal Sites. All pages belong to one of two types of page sets: public pages and private pages. By default, anyone can access public pages, even non-logged in users (guests). Only users who are members of the Site that owns the pages can access private pages. This means the private pages of an organization's Site are viewable only by Site members and members of the organization.

From *Site Pages* you can do several things:

1. Click the + button in the top right corner to add a new page.
2. Click options icons manage page or page set settings.

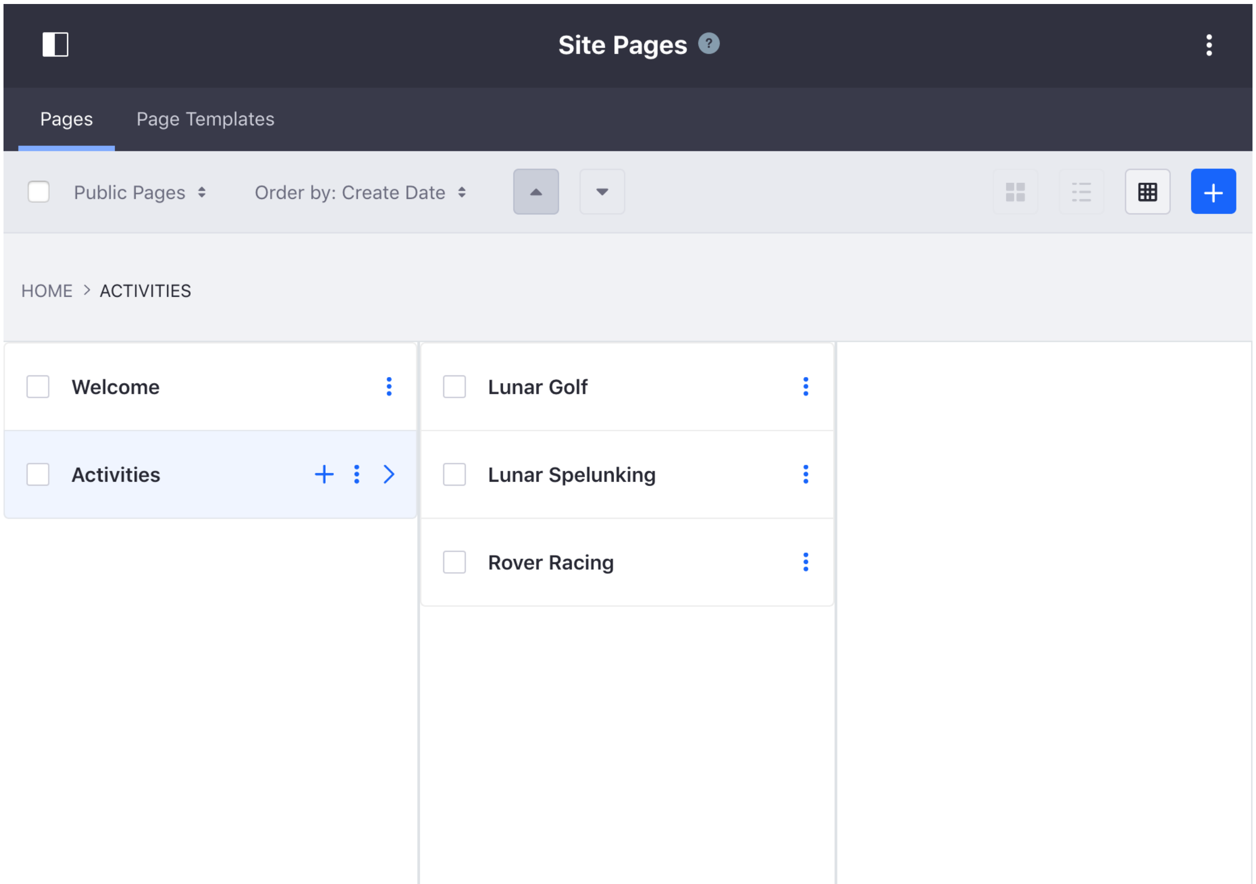

Figure 6.1: The Sites Pages page allows you to edit your Site pages as a whole.

3. Create child pages by clicking the + button next to an existing page.

Adding a child page creates child pages in the hierarchy below the page you've selected. You can nest pages as deep as you like.

Note: You're not forced to define the page hierarchy in a page's friendly URL. Therefore, child page friendly URLs are not required to include their parent page. For example, a parent page named Parent and its child page named Child could have the URLs *SITE_URL/parent* and *SITE_URL/child*, respectively. The default friendly URL given to a page is based only on the page name and not the hierarchy. If you wish to modify a generated friendly URL, you can do so by following the Friendly URL configuration section.

Once you've clicked the + icon to add a page, you're asked to select the type of page you are creating. There are two top options followed by other page types:

Widget: Creates a page with a layout template that defines a number of rows and columns for adding widgets to your page.

Content: Creates a Content Page with inline editing based on Fragments.

Below those you have other options:

Full Page Application: Creates a page that displays a single full page application.

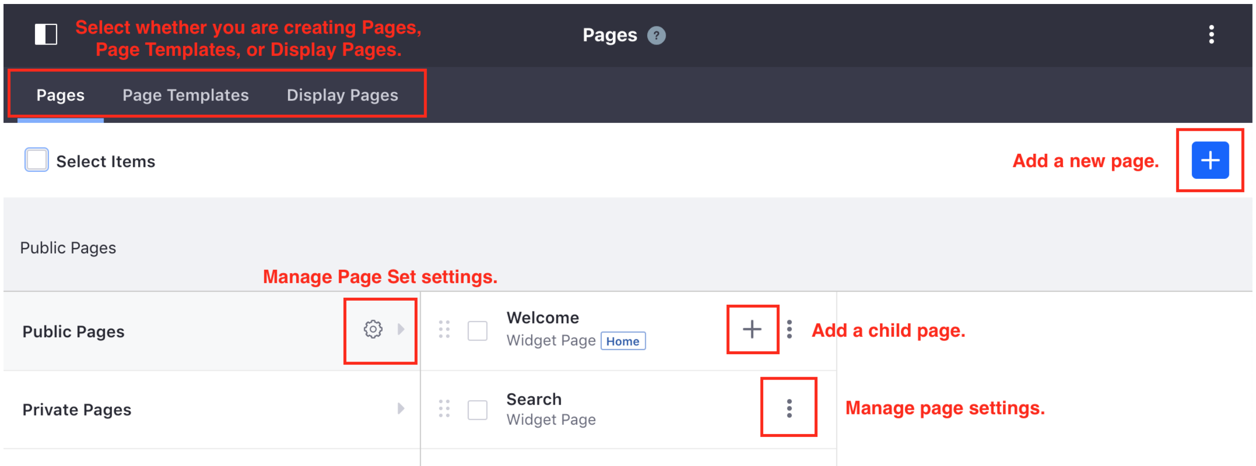

Figure 6.2: Understanding the options on Site Pages.

Page Set: Creates a container for subpages that is not actually a page itself.

Link to a Page of this Site: Links to a page within the same Site. This is often used to make a page available in multiple parts of a Sites hierarchy.

Panel: A page containing any number of applications as selected by an administrator, but only one is displayed at a time. Users select the portlet they want to use from a menu on the left side of the page, and the selected portlet takes up the entire page.

Embedded: Displays content from another website inside your instance. An administrator can set a URL from the page management interface and that page appears in the context and within the navigation of your Liferay instance. To embed an external website, you must provide the protocol in the URL (e.g. `https://www.liferay.com/`).

Link to URL: Creates a link to any URL. This could be an external page or a link across Sites in the same Liferay instance.

To the left, under Collections, you can choose to view the basic page types or a collection of page templates. By default, only *Global Templates* appears, but additional collections you create appear here as well.

After you've added a page, it may be difficult to track what kind of page you're currently viewing. The page type appears at the top of the page to help you determine the administration options you have and where you need to go to configure the page.

Now that you know the basics of adding pages, you can start working on the Lunar Resort Site. If you're not currently on the right Site, navigate to Site Administration in the Menu, select the compass icon next to the current Site name, and select the Site you wish to edit.

If you must ever modify the page you've created for your Site, select *Configure* from the Options menu for the page from *Site Pages*. When configuring a specific page, you have more options than when you were creating a new page. You can also read Configuring Sites.

There are also configuration options that are only available for individual pages or page groups only. You'll learn about options available for both instances.

Next, you'll look at creating the main page types you'll use in Liferay.

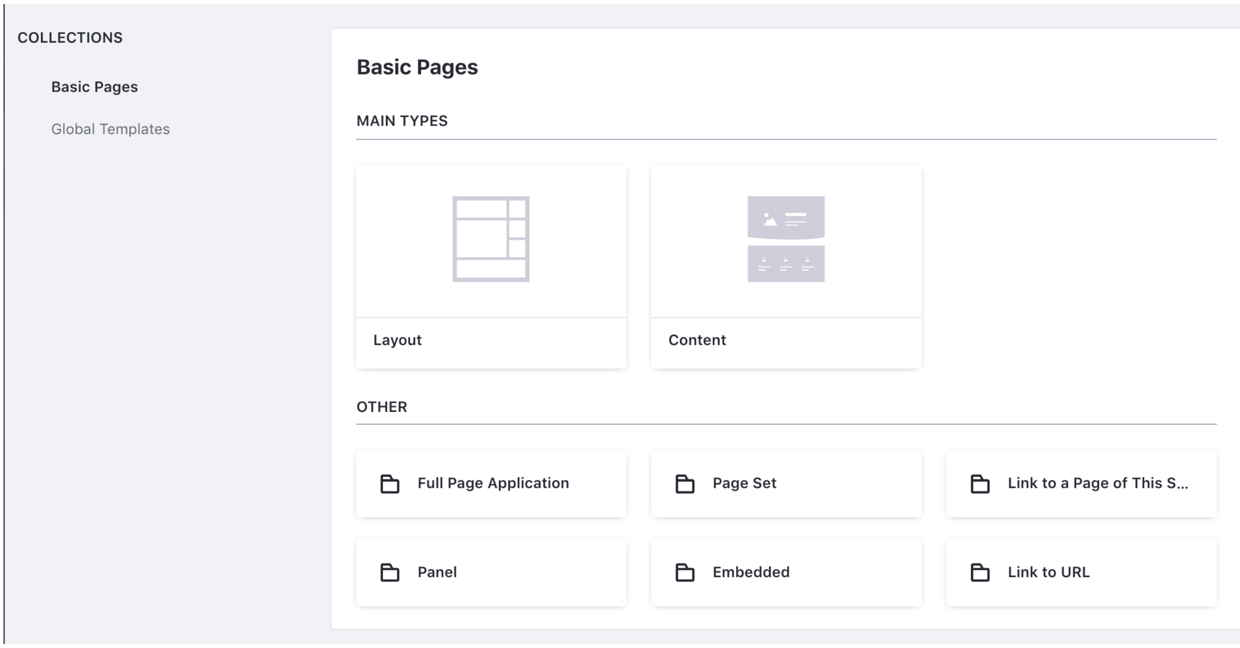

Figure 6.3: You must select a page type when adding pages.

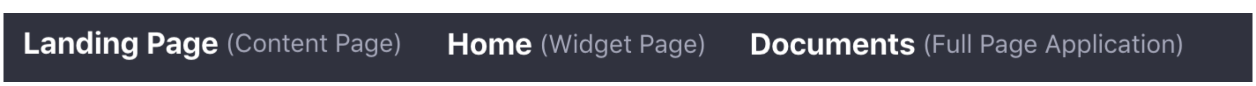

Figure 6.4: Here are three different page with three different types as they as displayed in the heading.

Chapter 7

Page Types and Templates

Each page type has unique features that fit together to create compelling sites. Here you'll look at each type in more detail, with a focus on the three most common page types: Layout Pages, Content Pages, and Full Page Applications. Since Page Templates are integral creating pages, you'll also learn about Page Templates and how to create and use Page Templates for different kinds of pages.

CHAPTER 8

USING CONTENT PAGES

Content Pages are new in 7.0. They provide a way to create pages with inline content. Traditionally in Liferay, all content has lived outside the page, to be added to any page as needed. With Content Pages, you can create a design on a page and then edit text or images directly within the page without any abstraction between the page and the content.

!PVideo Thumbnail

The process starts with a design. After the design is completed, a developer creates Fragments based on those designs. Widgets can be embedded in Fragments. Fragments can then be converted into Page Templates. Developers can define areas of Fragments to have editable text or images, allowing changes during the final publication process for the page.

This means that Fragments can be created as completely static elements to be used to build a Content Page, or marketers and content creators can customize them, using the Fragment as a tool for creating beautiful content.

!VVideo Tutorial

8.1 Creating Page Fragments

To create Content Pages, you must first have some Fragments. Fragments are one of the building blocks that you can use to create rich content in Liferay. Fragments are intended to be created by developers as they are built using HTML, CSS, and JavaScript.

Creating and Managing Fragments

Start in Site Administration. You can find Fragments in the Content section.

1. Open the main menu.

2. Under *Site Administration*, make sure the Site where you want to work is selected.

3. In the *Build* section, select *Page Fragments*.

Fragments are organized in *Collections*. The main Page Fragments page shows available Collections, provides the option to Import and Export, and enables you to create Collections. You

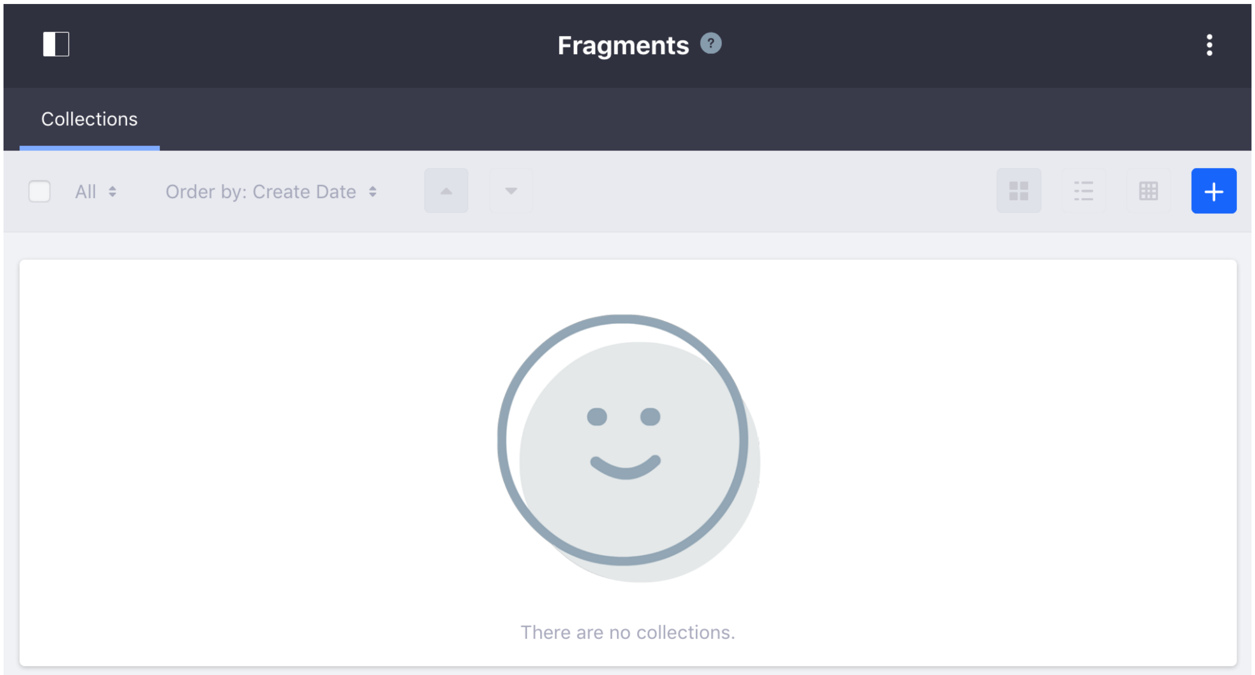

Figure 8.1: Here is the Page Fragments page with no Fragments or Collections created.

can also manage the organization and display of Fragments and Collections once you have them created. To create a Fragment, you must first create a Collection.

1. Click *New* → *Collection* to add a Collection.

2. Give the Collection a *Name* and *Description* and click *Save*.

Collections help you organize Fragments, and can be used to differentiate between different types of Fragments or Fragments used by different groups or departments. Next you want to create a Fragment inside the Collection you created.

1. Click on the Collection you created.

2. Click the *Add* icon ( ) to create a Fragment.

3. Give it a *Name* and click *Submit*.

Now you're looking at the Fragment development environment. Each pane in the editor has a different function:

- The top left pane is for entering HTML.
- The top right pane is for entering CSS.
- The bottom left pane is for entering JavaScript.
- The bottom right pane provides a live preview as you work in the other panes.

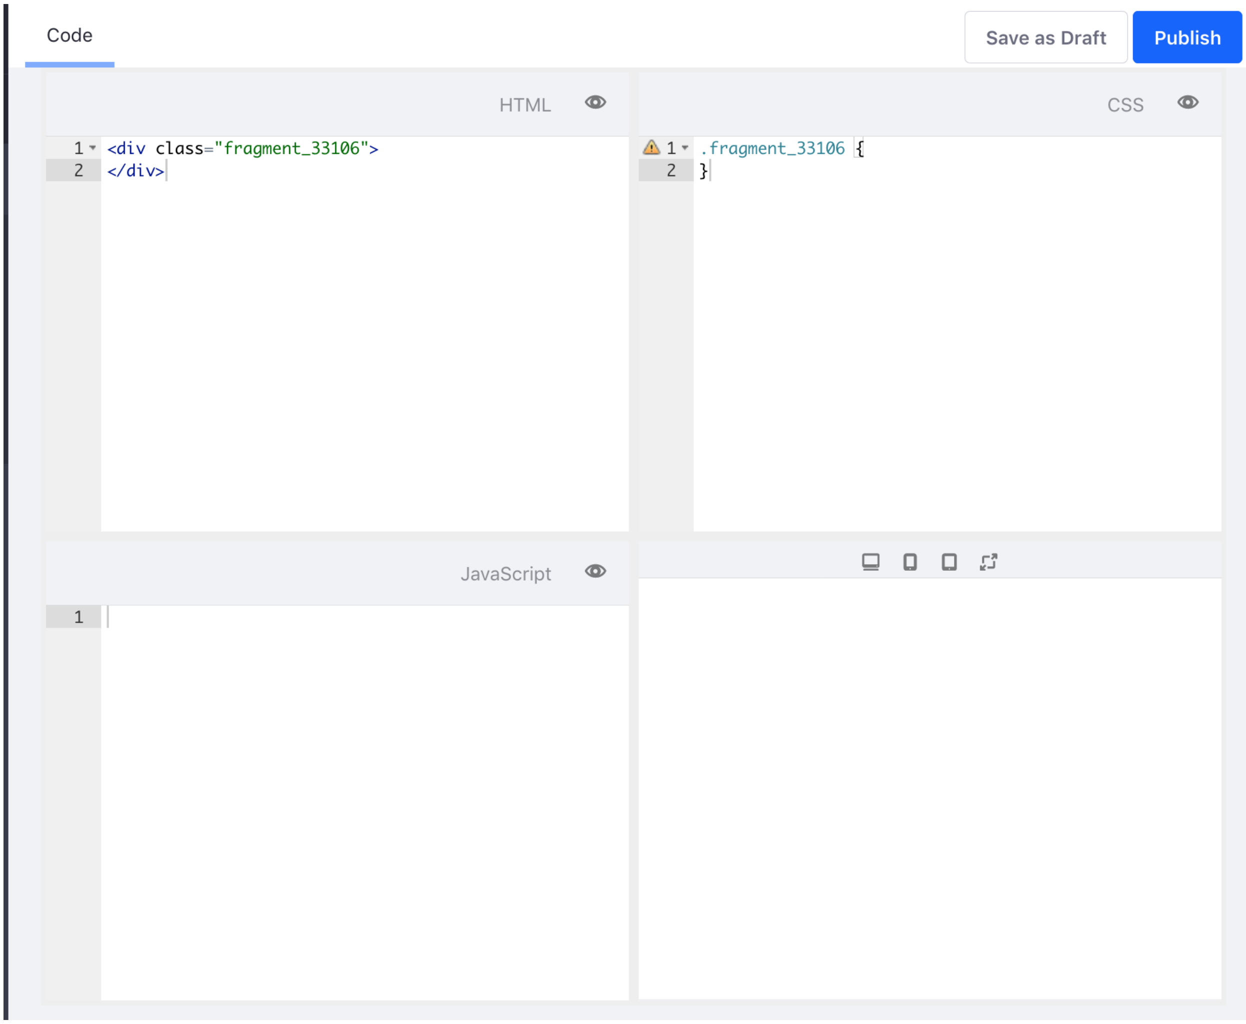

Figure 8.2: The Fragments editor provides an environment for creating all the parts of a Fragment.

In addition to standard HTML, CSS, and JavaScript, developers can also embed widgets and provide fields with editable text and images. The text and images can be edited during the final Content Page publication process. Fragment development is covered in depth in Developing Fragments.

After some Fragments have been created and published, you can start creating Page Templates to combine Fragments into pages.

8.2 Building Content Pages from Fragments

After Page Fragment collections are published, you are ready to create Page Templates. A Page Template is composed of some number of Fragments—one or fifty; it doesn't matter.

Creating a Page Template

You create Page Templates in the *Pages* page in Site Administration.

1. From Site Administration for your Site, go to *Build* → *Pages*.
2. Select the Page Templates tab. Like Fragments, Content Page Templates must be created in Collections. Your Collections appear on the Page Templates page.
3. Click the *New* button to add a Collection.
4. *Name* your Collection, provide a *Description*, and click *Save*.
5. Click on the new Collection.
6. Click the *Add* icon ( ) and select *Content Page Template* from inside the Collection.
7. Set the *Name* and click *Submit*.

Now you're on the Page Template creation page. Fragments are added by selecting a Collection from the *Fragments* tab on the right and adding them to the page. You can add multiple Fragments from different collections, and you can add the same Fragment to the page multiple times. To see the Fragments that comprise the current page, you can check the *Added* tab. The template is automatically saved as you work, but you must click *Publish* to make it available for use.

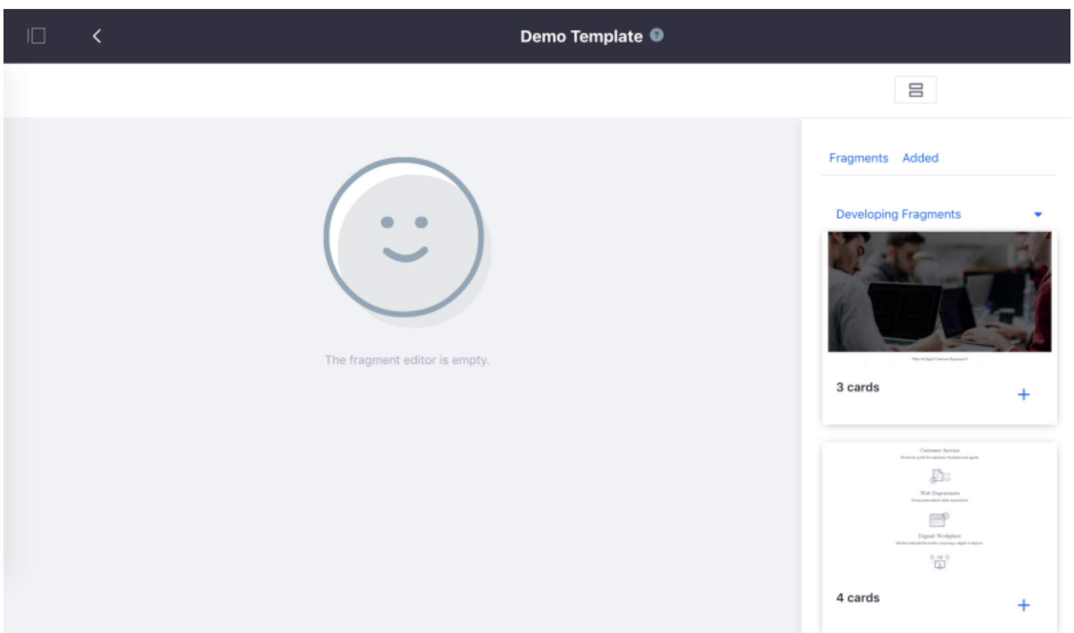

Figure 8.3: Drag Fragments to create a Page Template.

If the template contains editable text, you can click on the editable text area to change the text.

If the template contains an editable image, you can click on the image and replace it by uploading your own image or selecting one from Documents and Media.

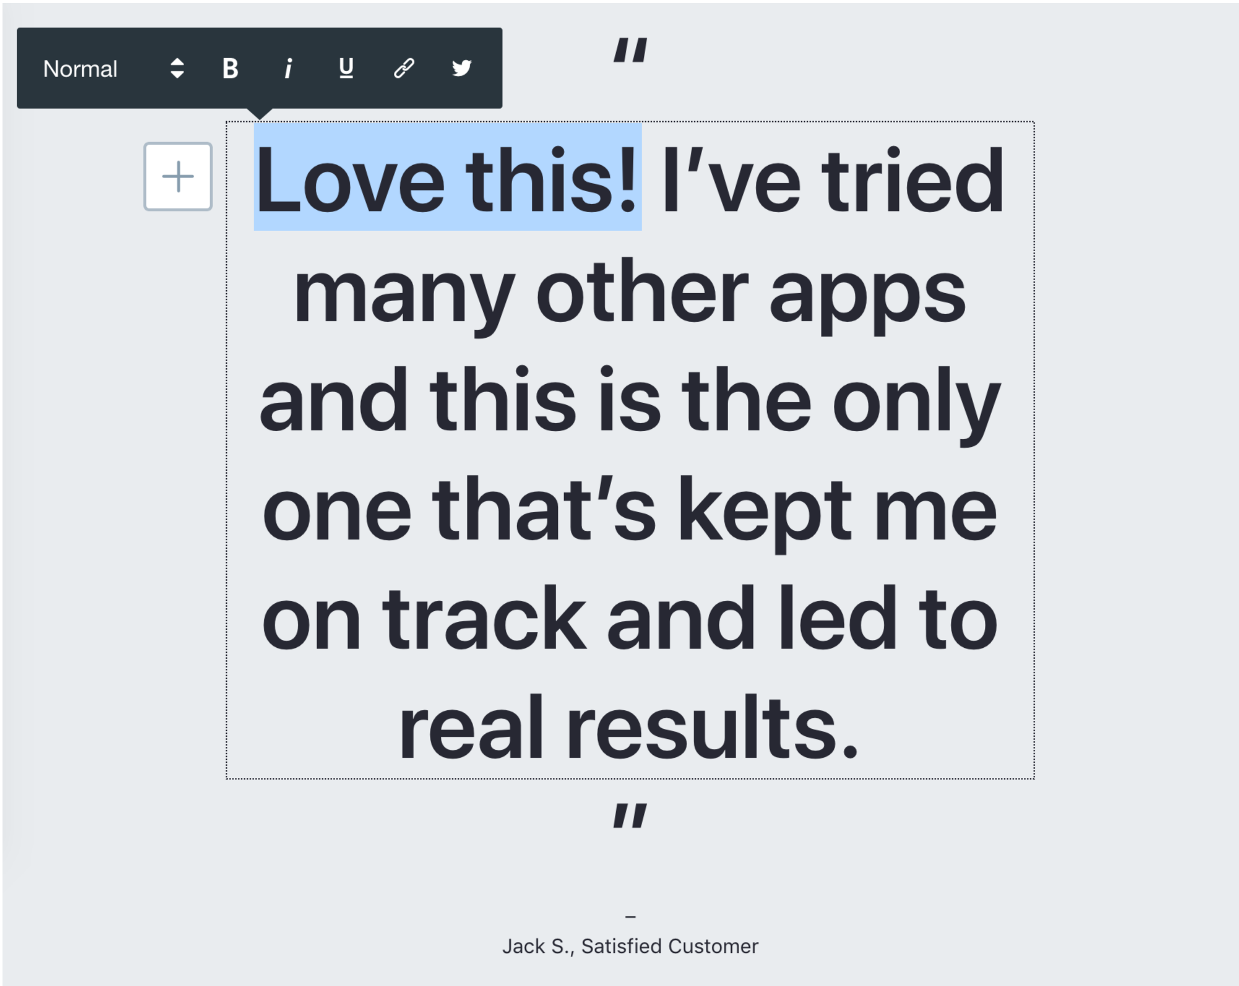

Figure 8.4: Editing text in-line lets you customize what's on the template.

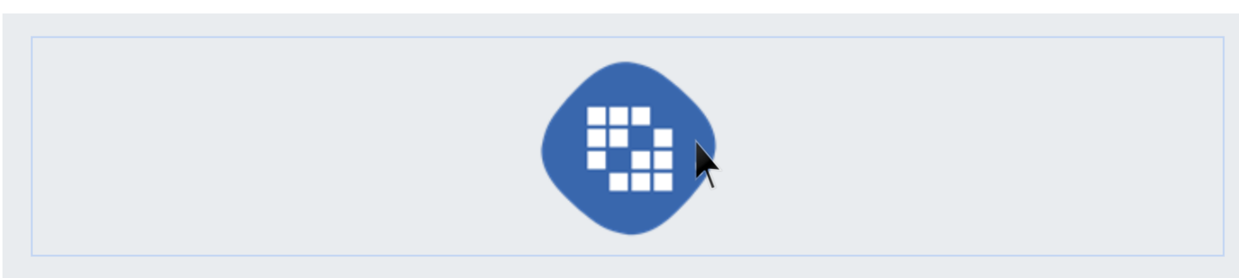

Figure 8.5: When you mouse over an editable image, a blue outline appears. You can replace it by clicking on it.

Note: While creating Page Templates, you can change editable Fragments. While creating the final Content Page, you can make changes to the Page Template. Changes you make to the Fragment for a Page Template are only reflected in that Page Template, and don't affect the Fragment itself. Likewise, when you apply edits to the text or images of a Content Page, those changes only exist on the current page and not on the Page Template itself.

Click on the back arrow at the top to stop editing the template.

Creating a Content Page

When you're finished creating a Page Template, you can use that template to create a Content Page. A Content Page is a page created from Fragments. Any attributes of the page that were defined by the developer as editable can be edited during the creation of the page or page template.

1. Go back to *Build* → *Pages*.
2. Click the *Add* icon ( ) for *Add Page* and select *Public Page*.
3. Select the Collection that contains the template you want to use from the *Collections* menu on the left.
4. Click on the template you want to use.
5. Enter a *Name*.
6. Click *Save*.

The page is published immediately, but you can edit the final page. You can change the chosen Fragments and edit any editable fields so that the final page is different from the Page Template.

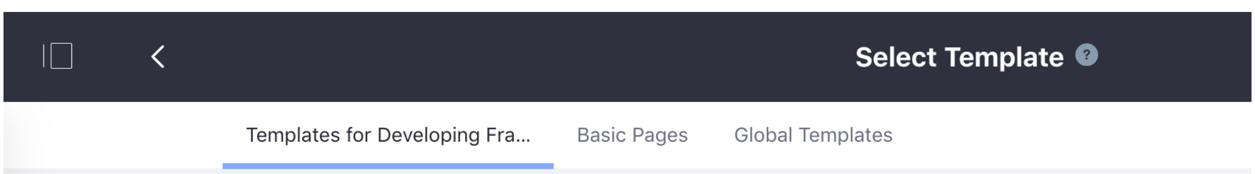

Figure 8.6: Selecting you page template.

Alternatively, you can create a Content Page without a template. From the page creation screen,

1. Select *Basic Pages* from under *Collections*.
2. Select *Content Page*.
3. Enter a *Name* and click *Save*.
4. Add Fragments to create your new page.

By default, your new page is added to the Navigation Menu and users can access the page you created.

Note: While portlets are rendered according to `render-weight` on Widget Pages, that is not true for Content Pages. Portlets are rendered in the order they appear on the page on Content Pages (i.e. left to right, top to bottom).

Propagation of Page Fragment Changes

If you make an update to a Page Fragment or Content Page Templates it doesn't automatically propagate changes, but you can access the *Usages and Propagation* page to selectively propagate changes.

1. From the Site Configuration menu, go to *Build* → *Page Fragments*
2. Select the Collection containing the changed fragment.
3. Open the menu for the fragment and select *View Usages*.

The *Usages and Propagation* Page shows a list of every Page, Page Template, and Display Page that uses the selected Page Fragment. You can then selectively propagate fragment changes to any or all of the pages listed. You can use the various filters and selection options to apply updates to pages quickly.

To update a page or template,

1. Select the page or pages you want to update by checking the box next to the page name.
2. Click the *Propagate* icon (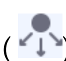)

After you propagate changes, visit any effected pages to verify there were no unexpected side effects of the changes. Changes to existing `editable` fields are not propagated since this overwrites content currently in content pages. To force propagation to content in an `editable` field, a developer must change the field ID. Any content created in that field will no longer display in the Content Page when the changes are propagated, but it will remain in the database and can be retrieved using the old ID.

Next you'll learn how to import and export Page Fragments.

8.3 Exporting and Importing Fragments

Often you'll want to reuse Fragments or re-purpose some of the code from Fragments in another Site or somewhere else altogether. Since all the content is plain text, copy/pasting your code would always be an option, but there's provides a much more elegant solution: exporting your Fragment Collections.

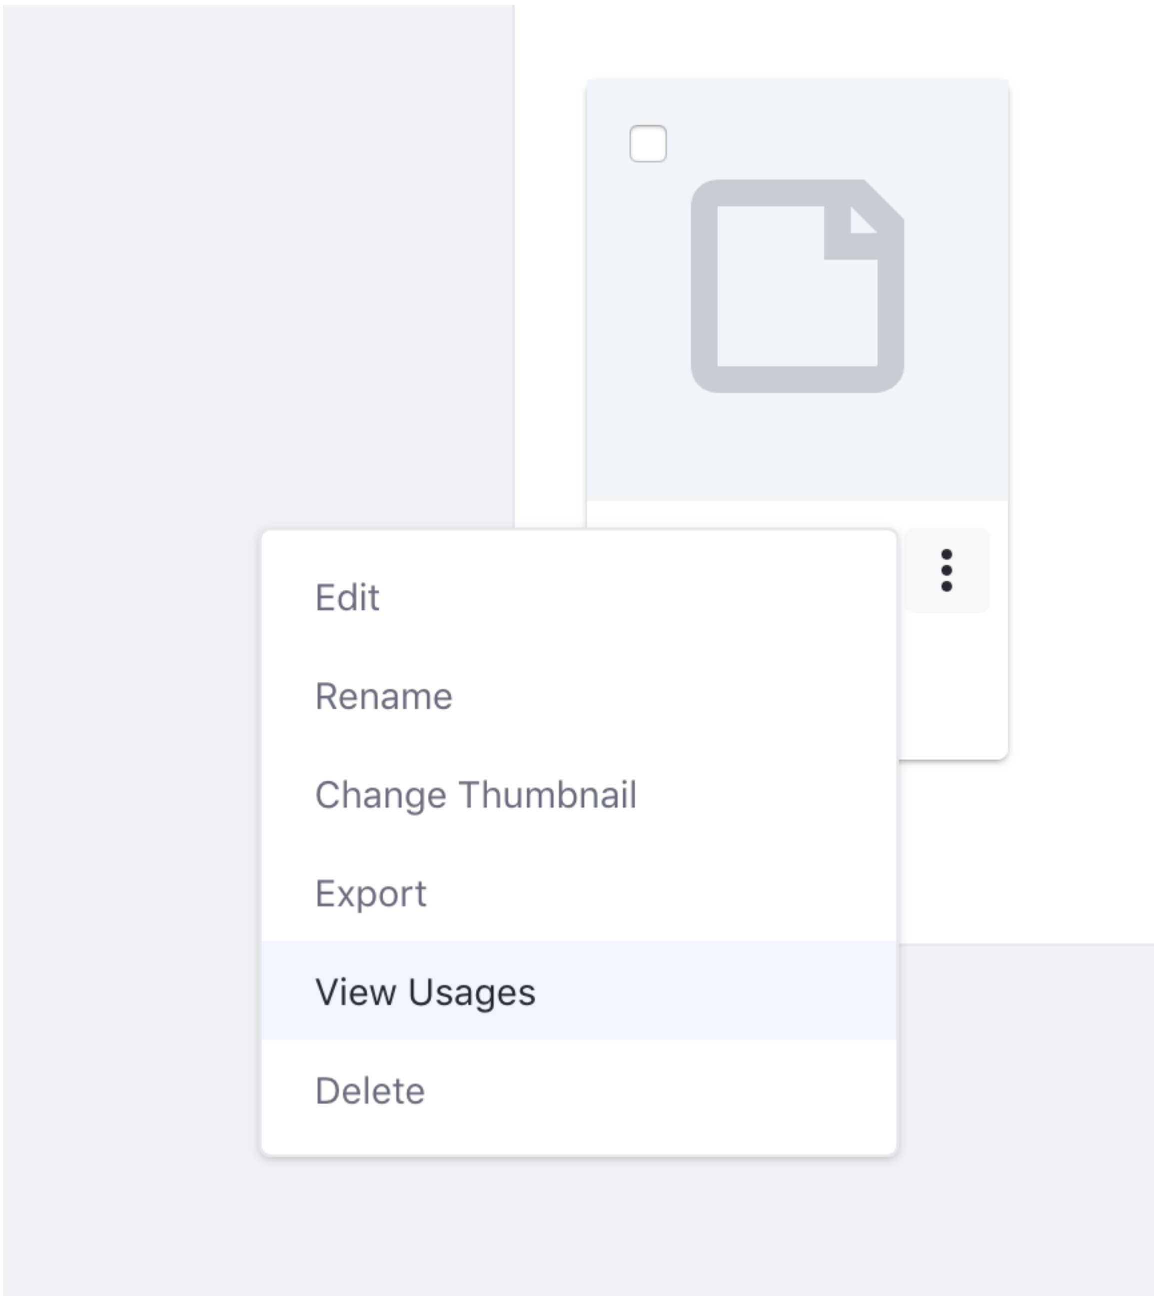

Figure 8.7: Select *View Usages* to propagate fragment changes.

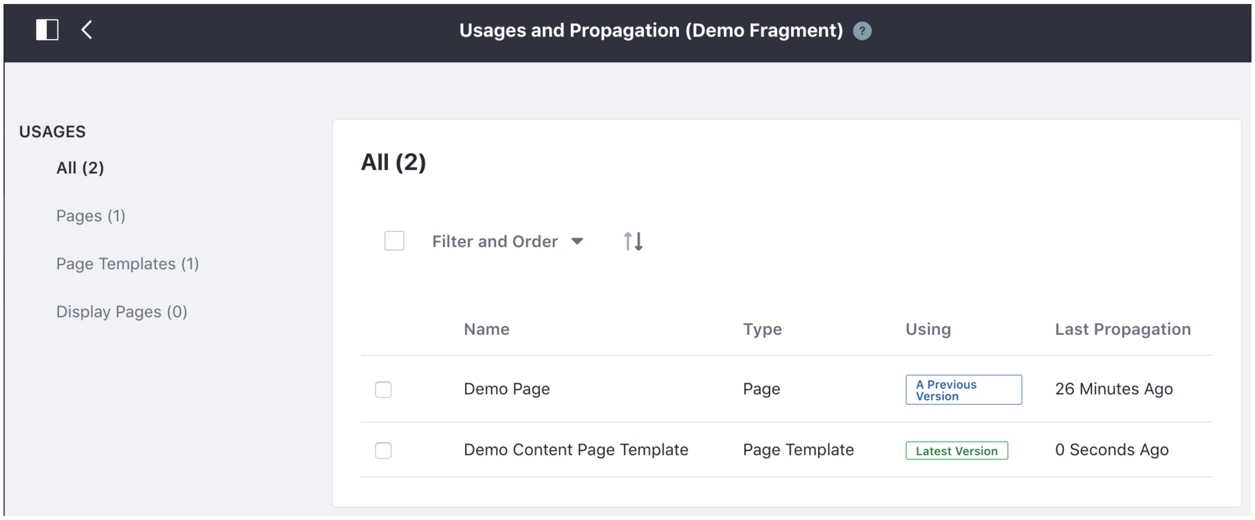

Figure 8.8: Viewing the Usages and Propagation page.

Exporting Fragments

There are two ways to export Fragments:

1. Export a single Collection.
2. Export some Fragments outside of a Collection.

To export a single Collection,

1. Go to *Site Administration* → *Build* → *Page Fragments*.
2. Next to *Collections* click *Actions* ( ) and select *Export*.
3. Select a Collection or multiple Collections to be exported. Each collection exports in a separate file.

Each Collection `.zip` contains all data for the Collection as well as the Fragments within it. To export individual Fragments,

1. Click on the Collection that you want to export the fragment from.
2. To export all Fragments in the Collection without exporting any Collection data, click *Actions* ( ) *Export* next to the Collection name. A `.zip` file is generated and downloaded automatically.
3. To export a single Fragment, click *Actions* ( ) *Export* next to the Fragment. A `.zip` file is generated and downloaded automatically.

When you export a single Fragment or a group of Fragments without a collection, you must have an existing Collection to import them into.

Now that you've done all this exporting, it's time to import it all back it.

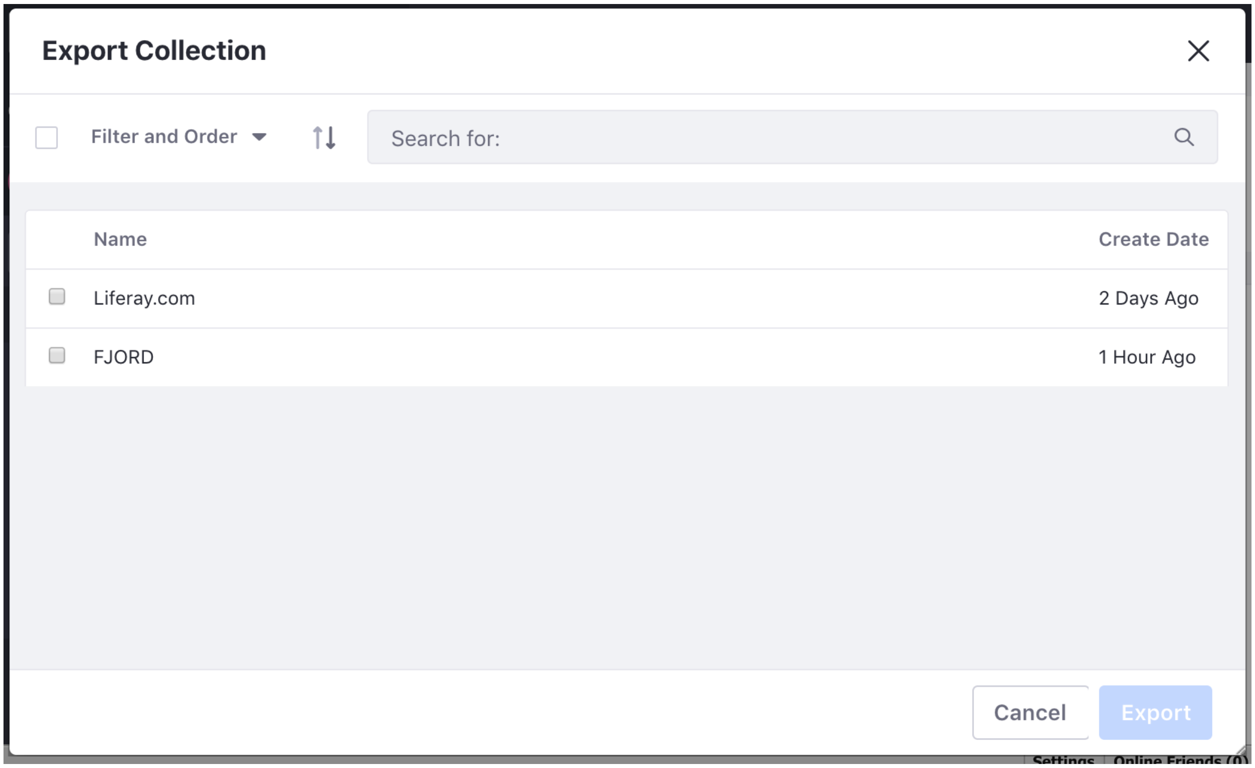

Figure 8.9: Select Collections to export.

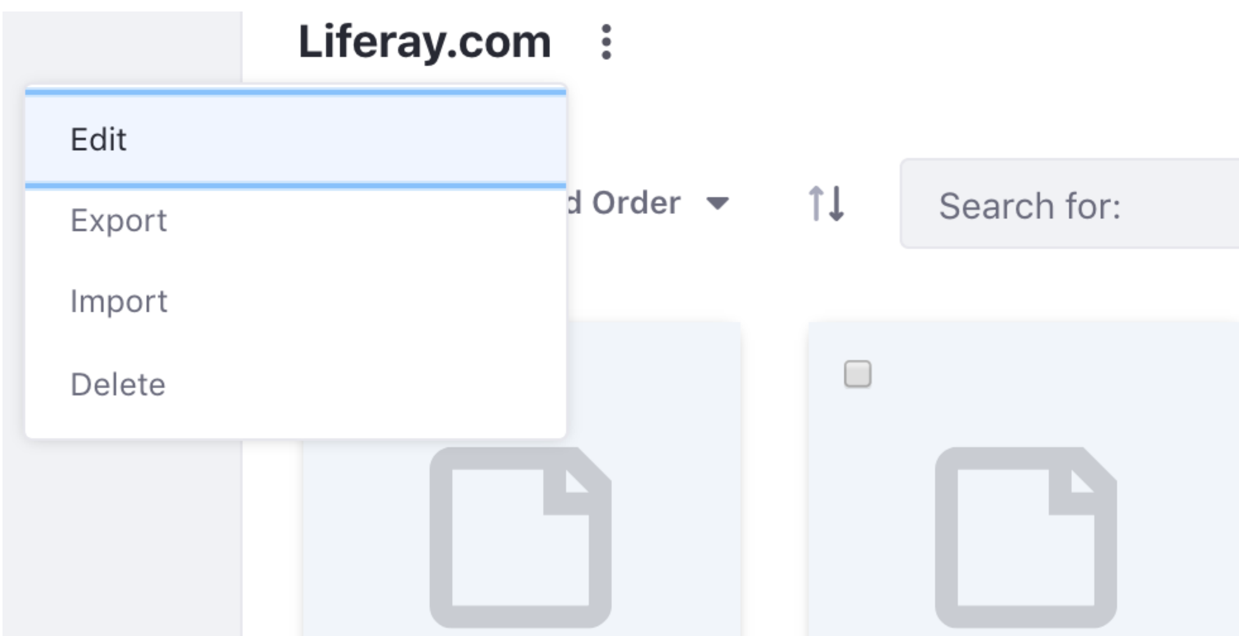

Figure 8.10: Exporting all of the Fragments in a Collection.

Importing

There are a few options for importing fragments, depending on how you exported them. You can import a Collection that was created in Liferay DXP, a Collection created using external tools, or any number of Page Fragments without a collection. Any time you import any Page Fragments they aren't available for use until you have gone to each imported fragment and approved it for use. This is to ensure that there are no errors in any imported fragments before they are added to a page.

See Developing a Fragment Using Desktop Tools for more information on creating and importing Fragments using other tools.

Importing Collections

You can import collections that were exported from Liferay DXP or that were created using other tools. To import a collection, follow these steps:

1. Go to *Site Administration* → *Build* → *Page Fragments*.
2. Next to *Collections*, click *Actions* (⋮) and select *Import*.

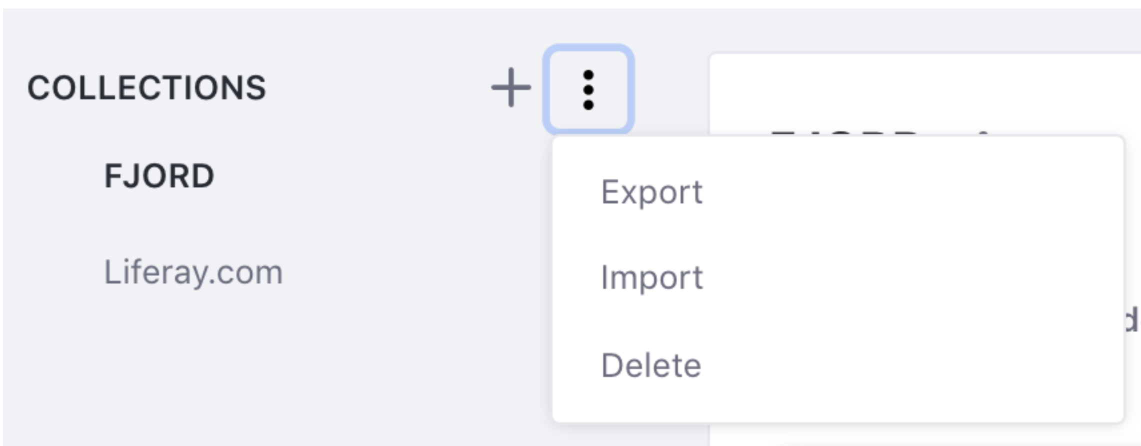

Figure 8.11: Importing and exporting Collections is accessed from a single menu.

3. On the next screen, click *Choose File* and select the file you want to import.
4. If you want to replace an existing collection, make sure the box is checked for *Overwrite Existing Files*.
5. Click *Import* and the collection is uploaded.

Importing Individual Page Fragments

You can also import a single Page Fragment or a set that was exported outside of a collection.

1. From the root level of the Fragments page, click on an existing Collection where you want to import the Fragments.
2. From inside the Collection click the *Actions* (⋮) button in the top right corner of the page.

3. Select *Import*.

4. Drag-and-drop or click *Select* to upload the Fragments `.zip`.

5. Click *Import*

The Fragments are imported into the Collection where you performed the import.

Export and Import let you use your preferred developer tools or share Fragment functionality elsewhere.

Chapter 9

Using Widget Pages

Widget Pages are composed of "widgets." A widget is any application that you can add to a page. Widget Pages are constructed by adding widgets to the page and filling them with content.

A widget could be a wiki display or a dynamic publishing tool like the Asset Publisher. The content you display with widgets could be long-form text or an image gallery, or anything in between. In this section, you'll learn to create widget pages and build content with them.

9.1 Creating Widget Pages

Widget Pages are the classic type of page in Liferay DXP. They're simple to create and fill with content or functionality. You create a blank page, define a layout, and then add widgets to the layout. Widgets can display content or provide some tool or function for Users.

Adding a Widget Page

When you first start Liferay DXP you get a widget page by default as your home page. To create a new widget page,

1. Go to *Site Administration* → *Build* → *Pages*.
2. Click the *Add* icon ( ) in the top right and select *Public Page* to add a new page.
3. Select *Basic Pages* if it is not selected by default.
4. Select the *Widget Page* type.
5. Name the page *Community* and leave the box checked to *Add this Page to the following Menus: Default*.
6. Click *Submit*.
7. On the next screen, you can select a Layout Template or manage other options. Leave the defaults and click *Save*.

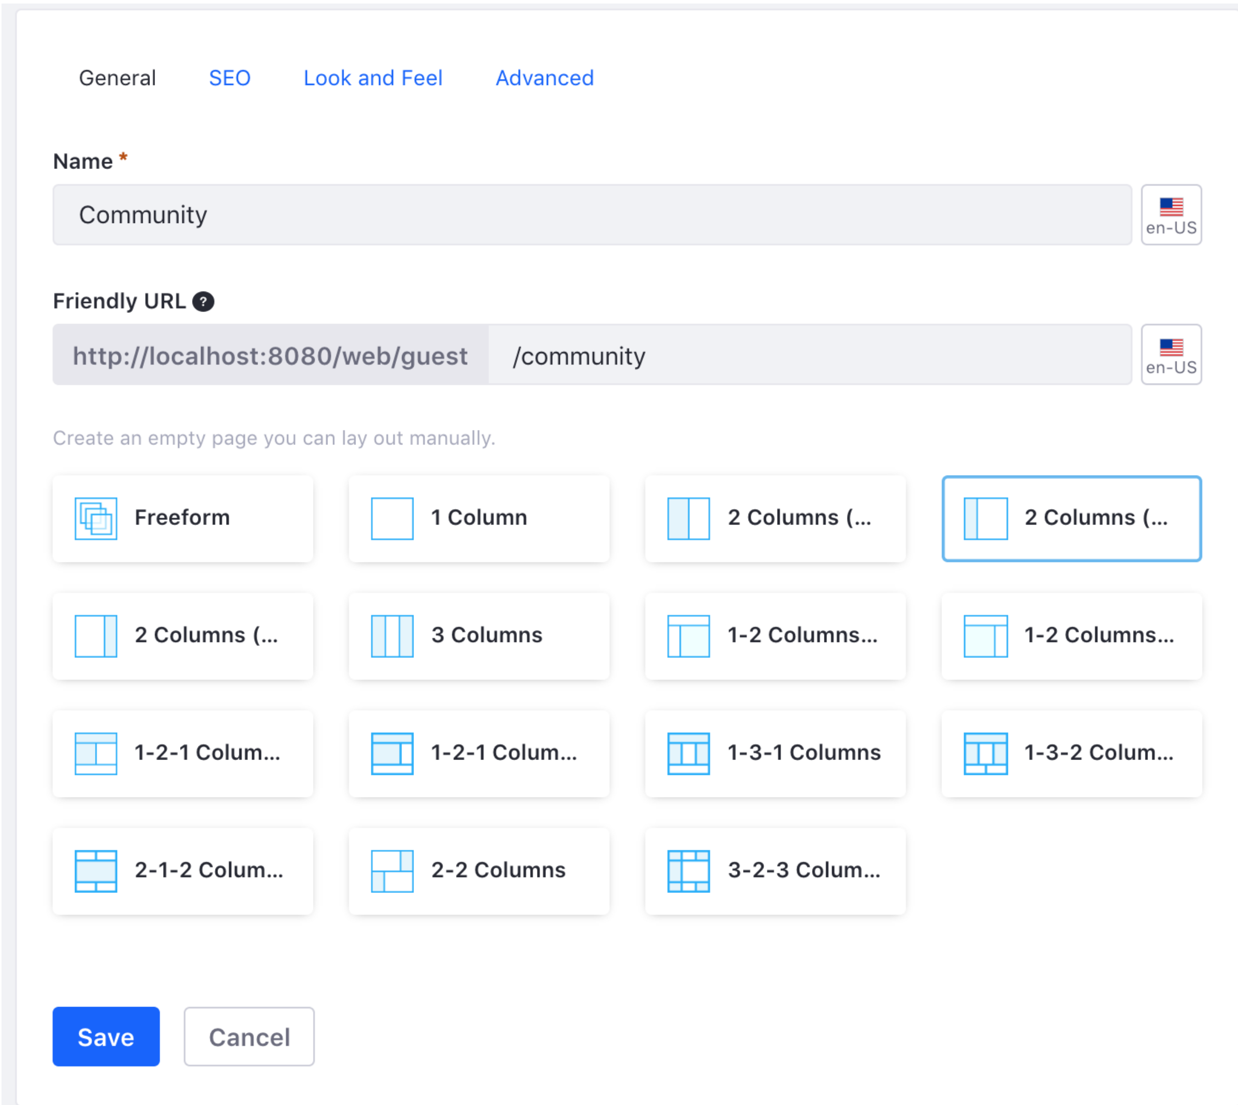

Figure 9.1: Create a page called *Community* with two columns.

Creating a page by default also adds it to any Navigation Menus that are configured to have new pages added to them. If you don't want a new page added to a specific Navigation Menu that is listed during page creation, uncheck the box for that menu.

Your new page is now added to the navigation.

1. Click the logo in the top left of the page to go back to your Site's front page. The page you just created appears in the main navigation.

2. Click on *Community* to go to the page.

Currently the page is empty. Next you'll add some widgets to give it functionality.

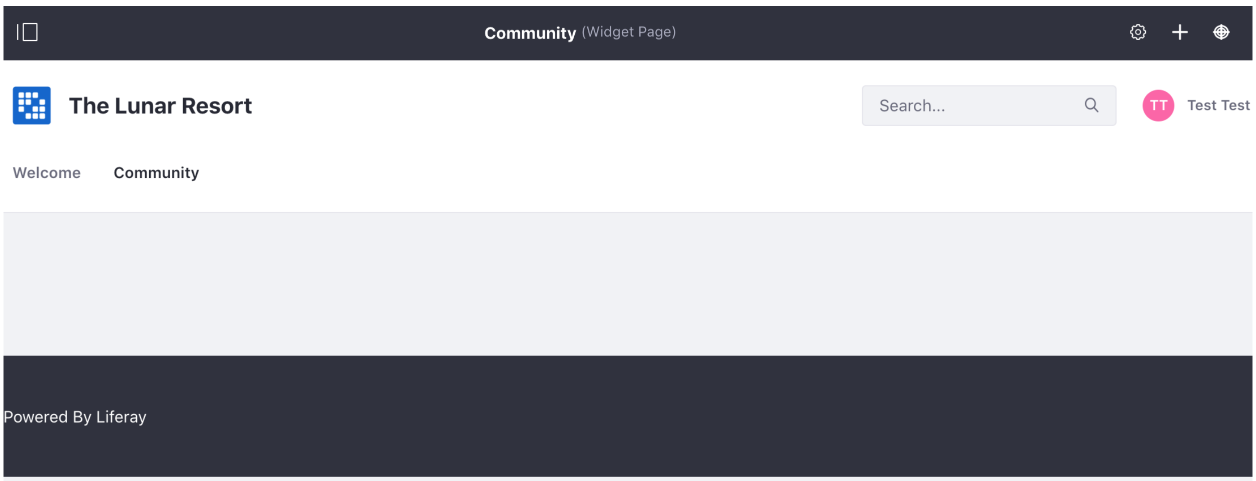

Figure 9.2: Your page has been added to the navigation automatically.

Adding Widgets to a Page

To add widget to a page, go to the page and click the *Add* button () from the top menu and select the *Widgets* tab. You can either browse through the categories of available widgets until you find the one you want, or you can search for widgets by name. Once you've found a widget, click the *Add* button to add it to the current page. Once there, you can drag it to a new position. Alternatively, you can drag the widget directly from the Widgets menu to a specific location on the page. Follow the steps below to add some Collaboration apps to the Lunar Resort Site.

1. From the top menu, select *Add* → *Widgets*.

2. In the menu that appears, expand the *Collaboration* category.

3. Drag the *Blogs Aggregator* widget onto the right column of your page.

4. Next, drag the *Wiki* app to the left column.

See how easy it is to add applications to your pages? You've added the Wiki app and Blogs Aggregator app to a page.

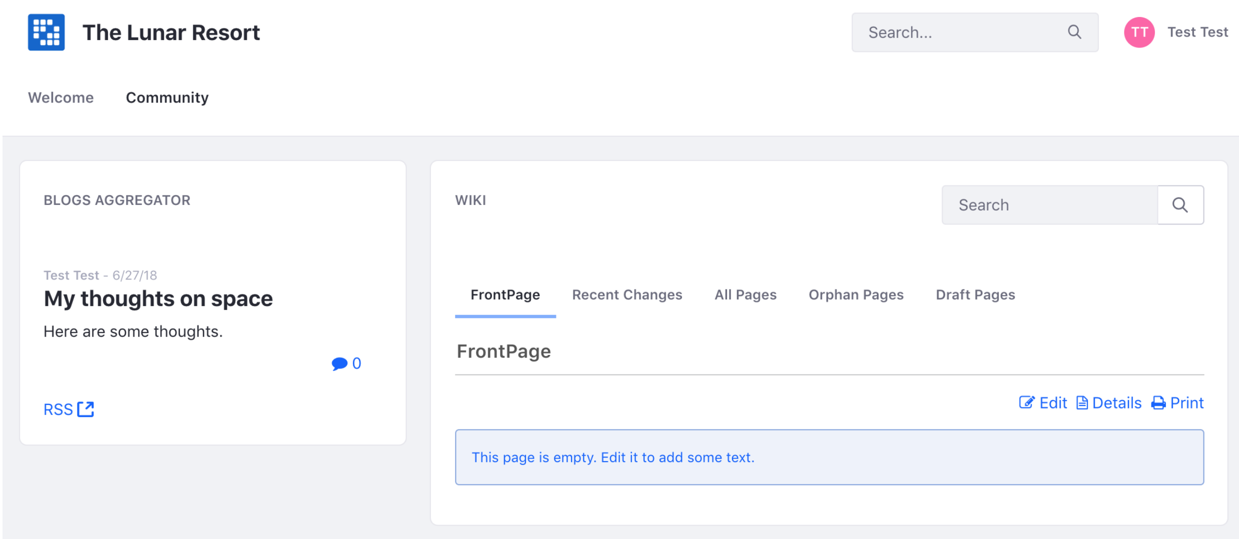

Figure 9.3: Your page layout options are virtually limitless with a slew of application and layout combinations.

If the default layout options provided aren't enough, you can create your own. For more information about developing custom layout templates, see the tutorial Layout Templates with the Liferay Theme Generator.

Note: Portlets are rendered according to `render-weight` on Widget Pages.

Next, you'll look at creating reusable templates for widget pages.

9.2 Creating Widget Pages from Templates

Page templates provide pre-configured pages to reuse. There are two types of page templates in 7.0: Widget Page templates consist of a portlet layout and configuration. Content Page templates are constructed from Fragments. You can read about Content Page Templates in this article.

Three sample layout page templates are installed by default:

- **Search:** Contains a search bar and configuration to display various facets.
- **Wiki:** Provides a page with three applications related to authoring a wiki.
- **Blog:** Provides a page with three applications related to blogging.

To add a new widget page template,

1. Go to *Build* → *Pages*.
2. Select the *Page Templates* tab.
3. Create a Collection named *Lunar Resort Templates*.

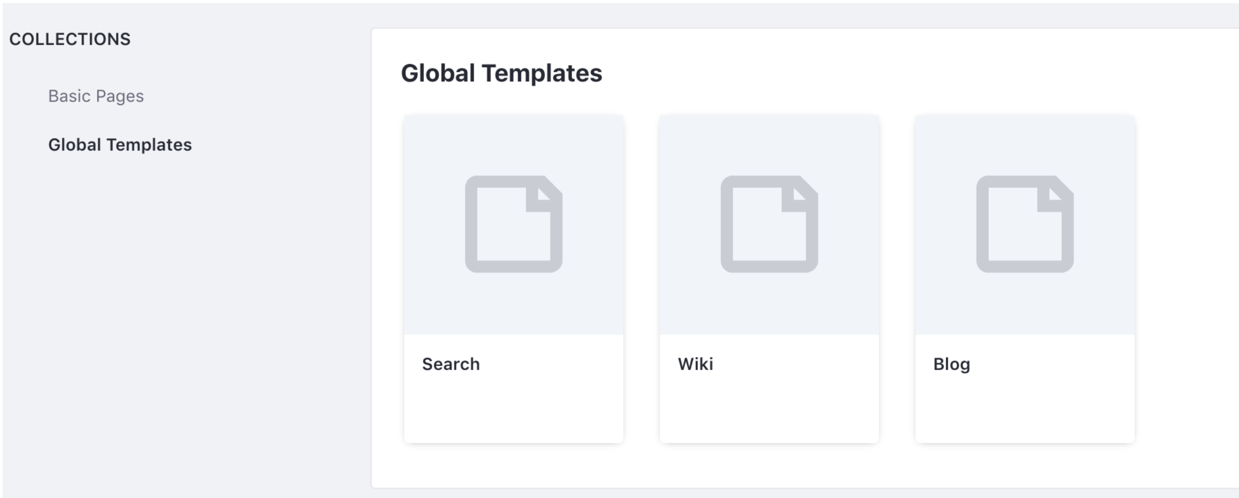

Figure 9.4: The Blog page template is already available for use along with the Search and Wiki page templates.

4. Click the *Add* icon ( ) and select *Widget Page Template.
5. Enter a *Name*.
6. Click *Save*.

The editing page for the template appears. You can add widgets to the page or access page configuration now. The changes you make are instantly applied to the template.

If you want to edit the template again,

1. Go back to the *Page Templates* tab.
2. Click the *Actions* icon ( ).
3. Click *Edit*.

Note that after a new page template has been created, the default permissions only allow the creator to use the page template. To give other users access to it,

1. Use the Actions menu for the template and select *Permissions*.
2. In the matrix of Roles and permissions, check the *View* permission for the Roles needed to see the page template in the list of available page templates when creating a new page.

If you want any user who can create a page to be able to use the page template, check the *View* permission for the *User* Role.

To use your template to create a new page,

1. Go to Site Administration and select the *Pages* option from the *Build* menu dropdown option.
2. Click the *Add* icon ( ).
3. Inside the Lunar Resort collection, select the page template that you created.

4. Enter the name of your page and click *Submit*.

Pages based on templates can inherit changes from the page template:

Figure 9.5: You can choose whether or not to inherit changes made to the page template.

By default, when a Site administrator creates pages based on a page template, future changes to the template are automatically propagated to those pages. Site administrators can disable this behavior by disabling the *Inherit Changes* selector. Occasionally, propagation for page templates fails due to unintended errors. To learn how to manage a failed page template propagation, visit the Propagating Changes from Site Templates to Sites tutorial.

If staging has been enabled, changes to the page template are automatically propagated to the staged page. These changes must still be approved before the page is published to live. For this reason, the automatic propagation of page template changes to the staged page cannot be turned off and the *Inherit Changes* selector does not appear.

You can read more about staging in the Staging Content for Publication tutorial.

Sharing Widget Page Templates

When importing pages to a new site or environment, you must also import templates associated with those pages. Generally templates are included automatically when an associated page is exported, but if not you can export the template collection separately so the page can be imported to the new environment. To export page templates,

1. Go to *Site Management* → *Publishing* → *Export*.
2. At the top right of the page, click the *Add button* (+) to create a new custom export process.
3. On the *New Custom Export* page, under the Content section, listed under Pages, you can click the *Change* link to choose which Collections and templates are being exported.
4. When you're done configuring the export, click *Export* and save the exported LAR file.
5. On the target environment, go to *Site Management* → *Publishing* → *Import* and click the *Add button* (+).
6. Upload the LAR with your template data. If the LAR contains additional content you don't want to import, you can deselect it.

Once the template has been imported, the page can be imported normally to your new environment. For more information on exporting/importing content, visit the Importing/Exporting Sites and Content article.

9.3 Using Full Page Applications

Full Page Applications are the ideal way to display a Message Board, Wiki, or other application that demands a full page.

Configuring the Page

Creating a Full Page Application starts just like creating any other type of page.

1. Go to *Site Administration* → *Build* → *Site Pages*.
2. Click the ( ) icon.
3. Give your page a *Name* and click *Submit*.

At this point you created the page, but it contains no content. If you visit the page, there is no way to add any content or widgets to the page. You must configure the page for it to function.

1. From *Site Pages* click the ( ) button and select *Configure*.

   On the next screen, you can change the page's *Name*, *Friendly URL*, and set the *Full Page Application* as well as access other page configuration options in the other tabs.

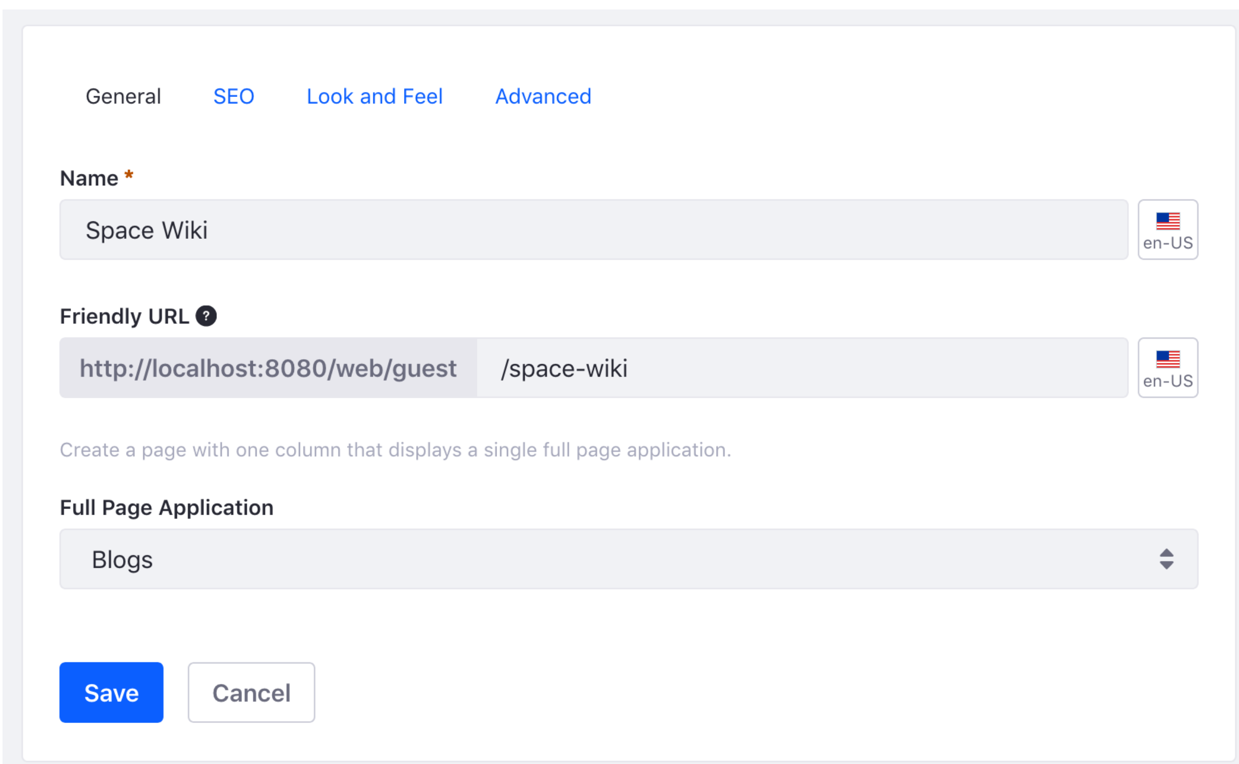

Figure 9.6: The Full Page Application configuration page.

2. Set the *Full Page Application* to *Wiki* and click *Save*.

   Out of the box, you can set the Blogs, Wiki, Media Gallery, Message Boards, RSS, Hello Soy Portlet, Documents and Media, or Dynamic Data Mapping Form to be the sole application for the page. Developers can make their applications Full Page Applications.

3. Click *Go to Site* in the Site Administration menu, and then click on your page.

Now the page is configured to display the Wiki and only the Wiki. No other widgets can be added to the page, and the Wiki app cannot be removed.

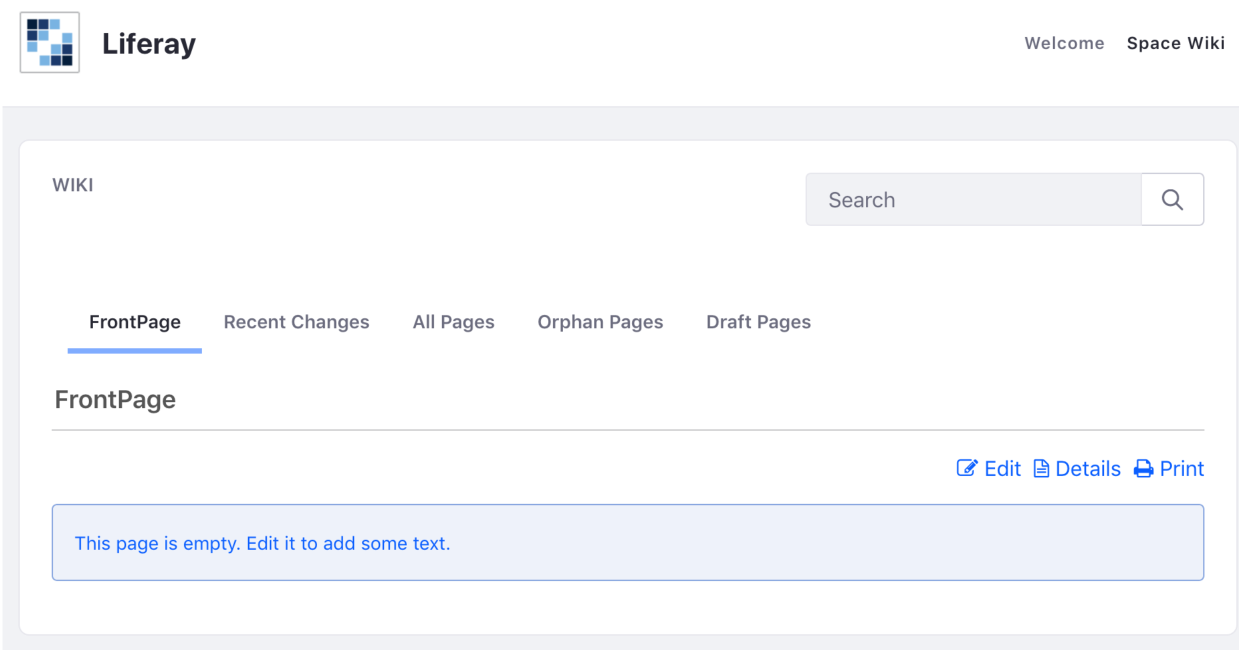

Figure 9.7: The Wiki displayed as a Full Page Application.

Note that all of the applications that can be added to the page are non-instanceable and the content of whichever application you select is based on the instance for that site. So if you already had data in your Wiki it appears on this page.

If you want to configure the application to be scoped to this specific page, you can configure that through the application's settings.

1. From the page, click the ( ) button for the Wiki and select *Configuration*.

2. From the *Wiki - Configuration* page, select the *Scope* tab.

3. Open the *Scope* menu and select *Space Wiki*.

Now the Wiki is scoped for this page, and doesn't share data with the Site or globally scoped Wiki.

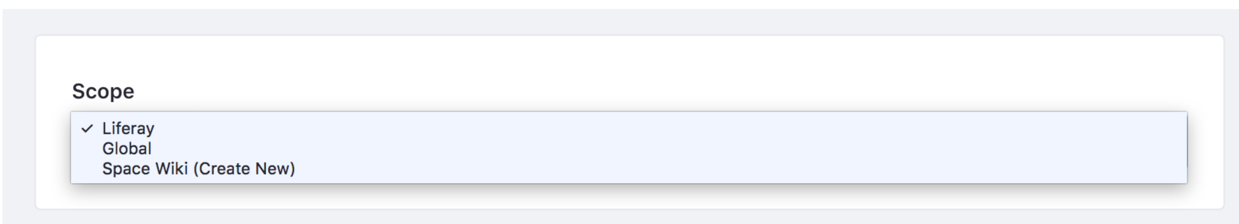

Figure 9.8: Configuring the scope.

CHAPTER 10

BUILDING A RESPONSIVE SITE

Now more than half of all page views in the world come from mobile devices like phones and tablets. That means that if your pages don't look good on mobile devices, your pages don't look good for more than half the people looking at them. Liferay DXP can provide the best experience possible no matter what device you're using.

Most of the heavy lifting for mobile optimization comes from your developers. Themes must be designed to be responsive. Web Content Structures and Templates and Page Fragments must be created with code that transitions gracefully to fit smaller screens on mobile devices. But there's work to do for marketers and administrators as well.

10.1 Built-in Mobile Support

Out of the box, there are several features that help make your pages look just as good and have the same functionality on mobile devices as they do on a desktop:

- Liferay Widgets and custom widgets that use Liferay's UI frameworks automatically scale to fit the screen size.

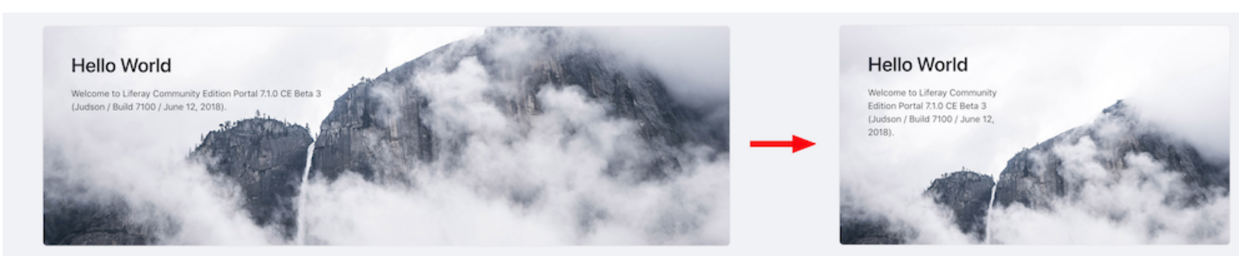

Figure 10.1: A widget adjusts its size.

- UI elements like the navigation and Product Menu automatically adjust to remain usable on smaller screens.

- When the screen width is low, Liferay combines columns so that all content remains legible.

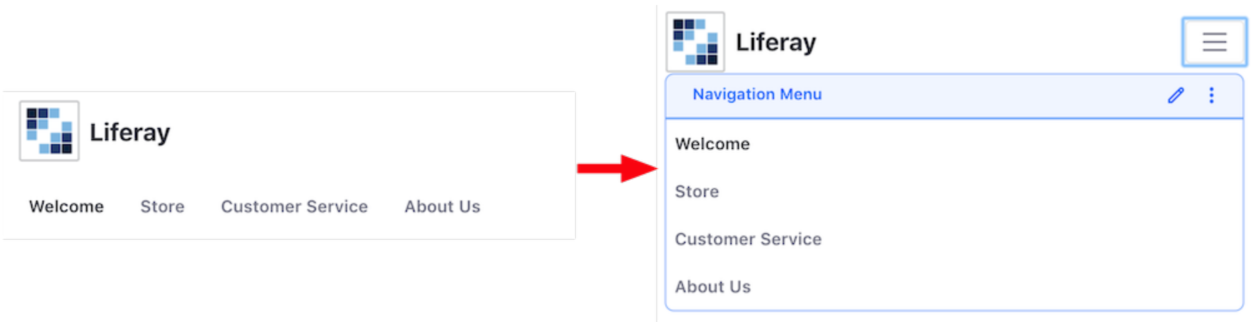

Figure 10.2: The main navigation adjusts its size.

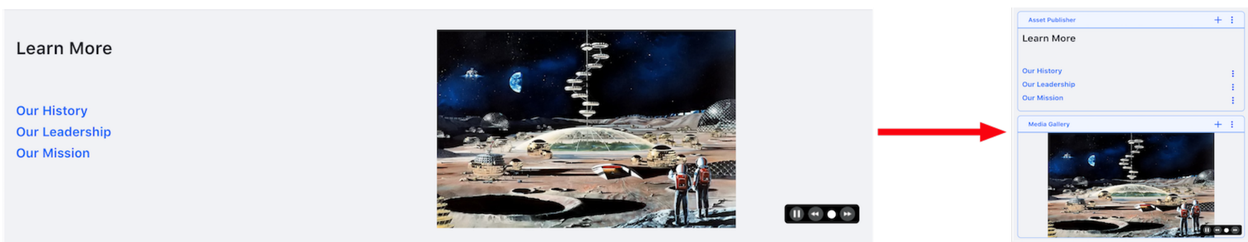

Figure 10.3: Columns combine.

- For web developers, Liferay's theme tools provide a number of tools to help ensure optimum mobile performance.

For most business users, this means that all you need to do to display pages on Mobile device is to create a page. However, you also have tools available to verify that everything displays as intended. The Device Simulator ( ) is a powerful tool that shows you how pages look on different devices.

Using the Device Simulator

When creating a page or reviewing a page before it is published, one of your most important tools is the Device Simulator found in the top right corner of every page. The simulator lets you view the current page in a number of resolutions based on different display types. There are three predefined options:

Desktop: Fixes the width to display the page at full size.

Tablet: Puts your page in a box as if it is being displayed on a tablet. It also activates some of Liferay's built-in mobile features.

Mobile: Puts your page in an even smaller box to demonstrate how the page looks to your average smartphone user.

There are also two options available to display

Autosize: Provides another way to view the default behavior where the page shrinks and grows based on the width of the browser window.

Custom: Lets you enter a specific size for testing and fixes the height and width of the display.

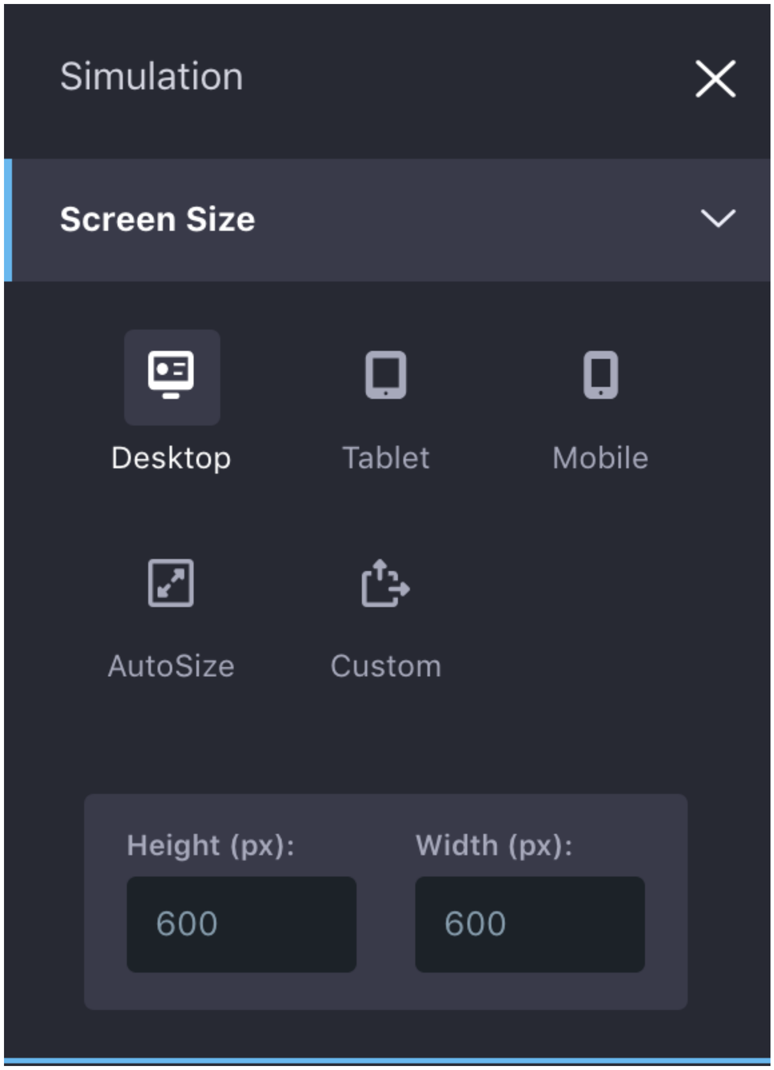

Figure 10.4: The Simulation panel defines multiple screen sizes.

Because modern mobile browsers are built on the same technology as desktop browsers, the behavior you see in the simulator should match the experience of users on mobile devices. In addition to making sure the basic layout looks good and that all functionality remains, it's also important to make sure that automatic features—like how columns are combined at lower resolutions—don't have unintended effects.

Designing Mobile Friendly Pages

Liferay DXP provides the tools you need, but building pages that provide a good experience across all kinds of devices still means working across all levels of web development and publishing. Theme developers must create themes that use Liferay's frameworks to scale content well across all kinds of displays. Designers must have multiple screen sizes in mind when designing pages. And before anything it published it must be thoroughly reviewed to make sure that it provides the best experience to all of your users.

Now that you've learned about Liferay's tools for making your website mobile friendly, let's look at your options for adapting to different types of mobile devices.

Chapter 11

Mobile Device Rules

With mobile device rules, you can alter what gets displayed based on the device being used to access Liferay DXP. For instance, you can configure the look and feel of pages accessed by smartphone or tablet users differently from those accessed by users on a computer.

Whole Sites or individual pages can be configured for mobile device families. A family describes a group of devices. You can set rules that describe a category of devices, such as all Android devices or all iOS tablets. You can define as many rules in a family as you need to classify all the devices for which you'd like to define actions. Families can be prioritized to determine which one applies to a given page request.

11.1 Creating Mobile Device Rules

To configure mobile device rules, you need a way to find out the characteristics of the device. While some of the characteristics are provided by the device, most are not. For this reason, there are databases that contain information about thousands of devices. These databases make it possible to learn every detail about a device from the device type. Liferay DXP's Mobile Device Rules can connect to device databases so that you can use their device characteristics in your rules.

Important: For the features described in this article to work, you must install the Liferay Mobile Device Detection (LMDD) app from Liferay Marketplace. This app provides the device detection database that's required to detect which mobile devices are accessing it. Note that if you're running Liferay DXP, you must install the lite version of LMDD before you can install the enterprise version. Click here for instructions on using Liferay Marketplace to find and install apps.

You can develop plugins that integrate with other device databases. Even if you don't have a device database, you can still set up mobile device rules. They won't, however, be effective until a database is deployed, because the portal won't have enough information about the devices being used to make page requests.

To access the Mobile Device Families administrative page,

1. Open the *Product Menu*.
2. Use the *Site Selector* ( ) to choose the Site that you want to define Mobile Device Rules for.

3. Select *Configuration* → *Mobile Device Families*.

You can also add families for all Sites by navigating to the Control Panel → *Sites* → *Global*. The Mobile Device Families administrative page displays a list of defined families and lets you add more. To add rules, you must first add a family.

1. Click *Add* button () to add a *New Device Family*.
2. Enter a *Name* and *Description*.
3. Click *Save*.
4. Click on the name of the Mobile Device Family to access the rules page.

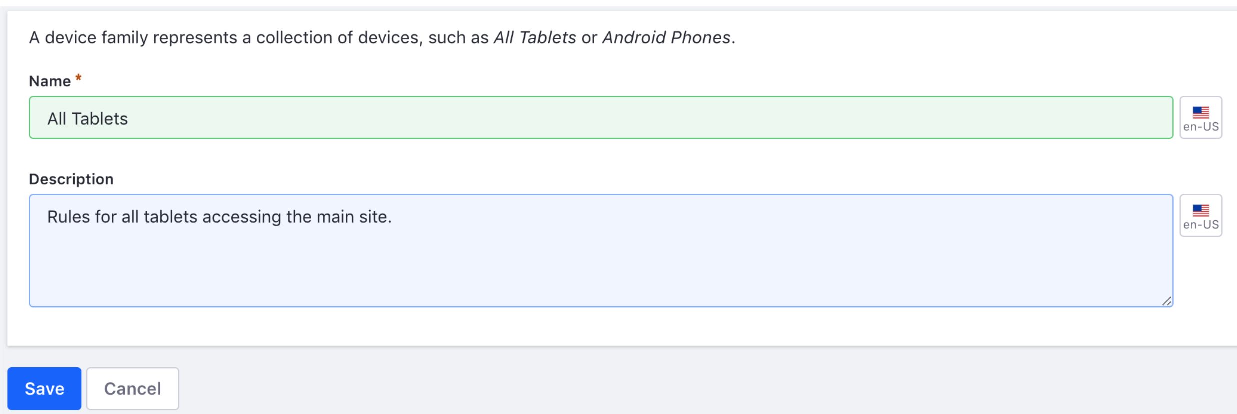

Figure 11.1: Create a Mobile Device Family so you can create rules.

The rules defined for a family, along with the priorities of the families selected for a particular Site or page, determine which family's actions are applied to a given request. From the New Classification Rule page for a specific rule set, you can add a rule by specifying an operating system, rule type, physical screen size, and screen resolution. Remember that you can add as many rules to a family as you need in order to classify the devices on which you'd like to take actions.

1. Click on the *Add* button () to add a new rule.
2. Enter a *Name* and *Description*.
3. Select the classifications you want for this rule from *Operating System and Type*, *Physical Screen Size*, and *Screen Resolution*.
4. Click *Save*.

You'll notice after saving the classification rule that it's characterized as a *Simple Rule*. Only Simple Rules are included with Liferay DXP, but the rules are designed to be extensible, and additional rule types can be added by your developers.

Once you've created some mobile device families and added some rules to them, you're ready to create some actions. The actions defined for a family determine what happens to a request when the device is detected and the family has been found to apply.

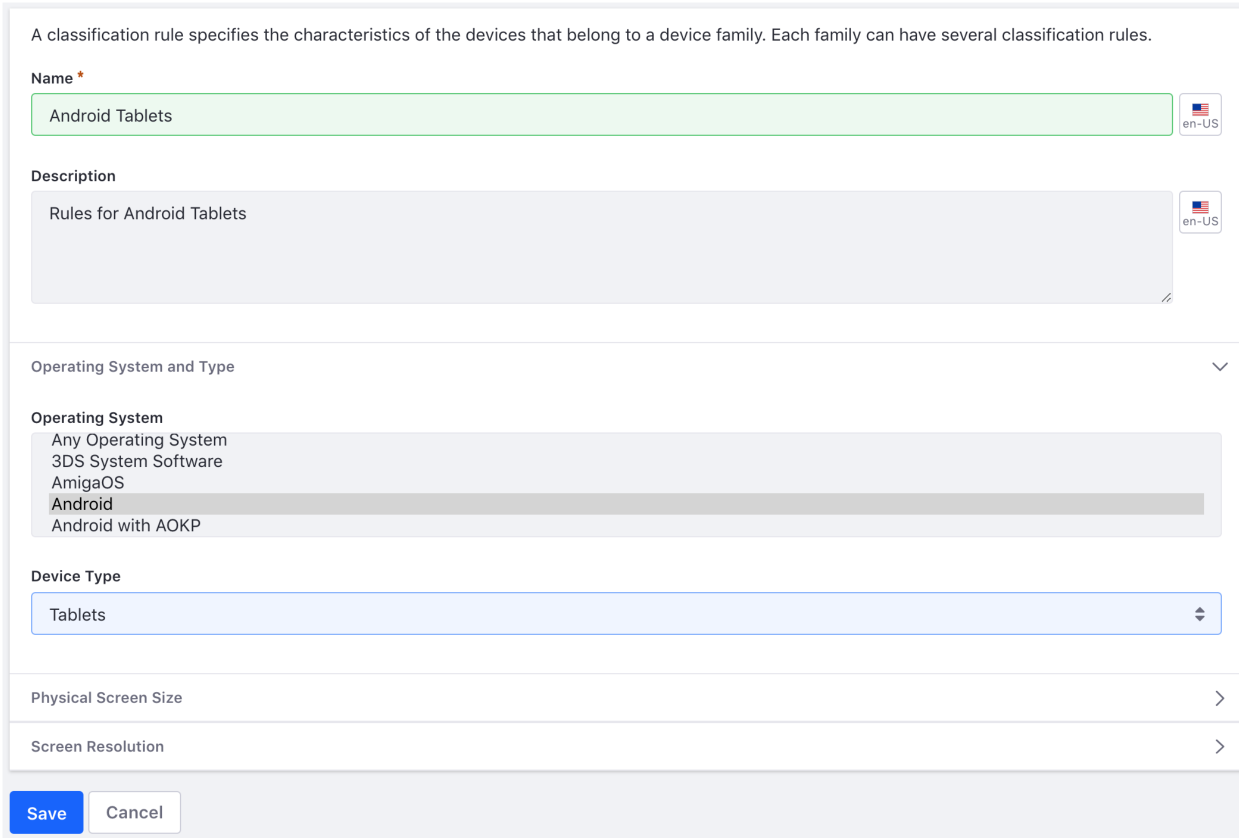

Figure 11.2: Select the operating system and device type for your rule.

Tip: The Audience Targeting application offers a *Device* rule that evaluates whether a User is accessing content using a particular device family. This rule is integrated with the Mobile Device Families app.

You can add families to a Site, individual page, or page set from their respective configuration pages. To do it for a Page Set:

1. Go to *Build* → *Pages* in your Site.
2. Click on *Configuration* () for the Public Pages.
3. Select the *Advanced* tab and open the *Mobile Device Rules* option in the bottom menu.
4. Click *Select* to open the list of families that can be applied.

From the same page, you can access the configuration for an individual page, or you can configure Mobile Device Rules for an entire Site from *Configuration* → *Site Settings*. You can select multiple families for a particular Site or page and order them by priority. The families are checked in decreasing order of priority: the actions defined by the first family that applies are executed.

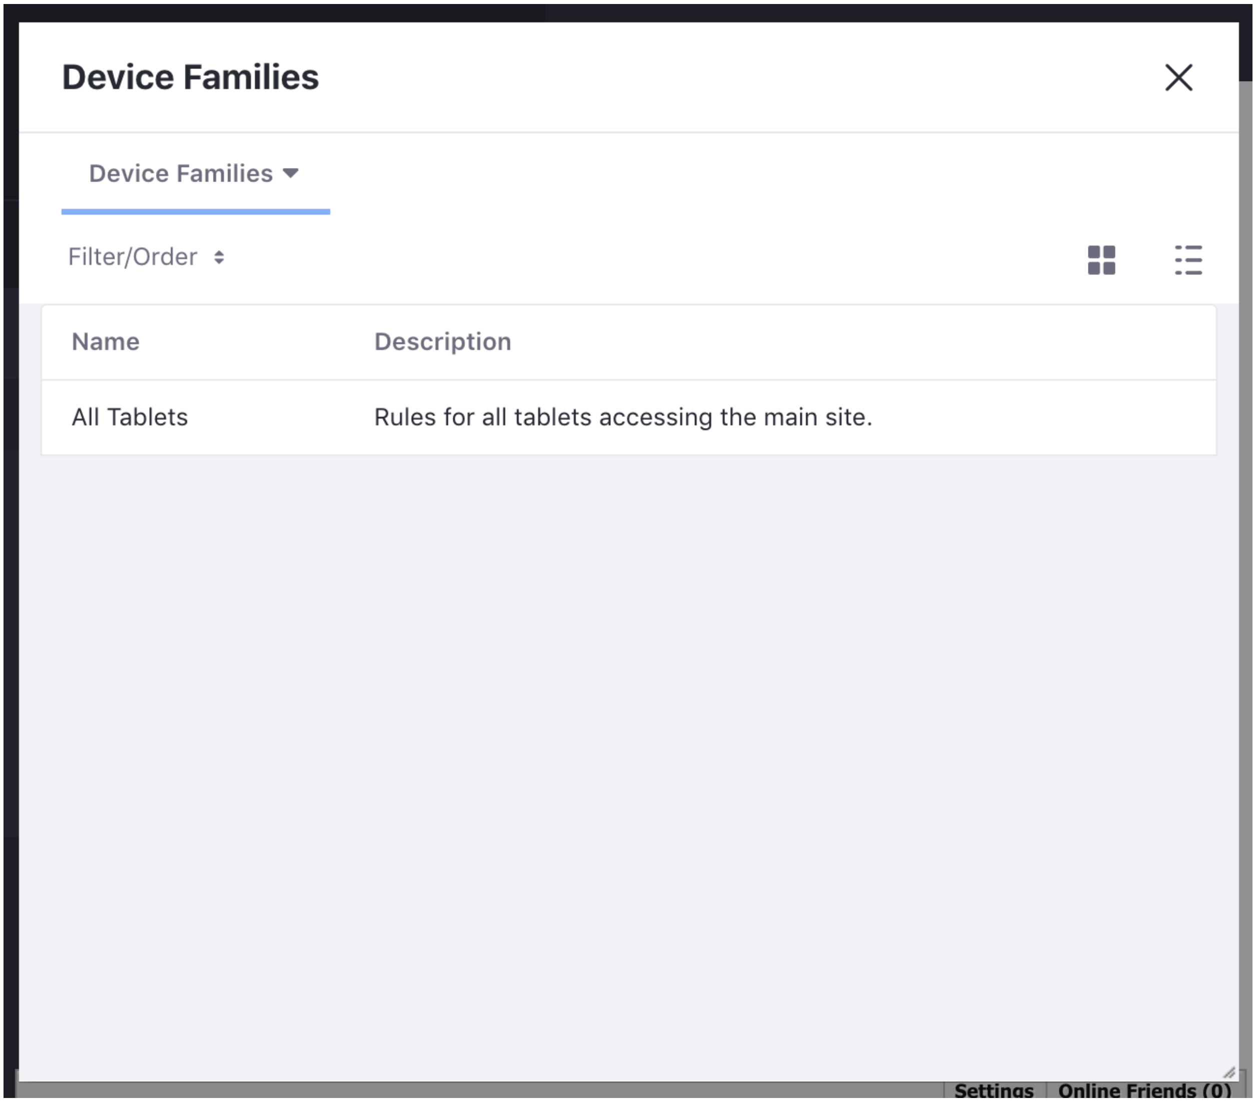

Figure 11.3: You can select a mobile device family to apply for a Site or page.

11.2 Mobile Device Actions

After you've created families and applied rules to define those families, you can associate specific actions that occur when a user visits that Site on a device.

To add actions to a selected rule group:

1. Go to the *Configuration* page for the page, page set, or Site where you have configured a device family.

2. In the *Mobile Device Rules* section, click *Actions* (⋮) → *Manage Actions* next to the device family that you wish to add an action for.

3. Click *Add Action*.

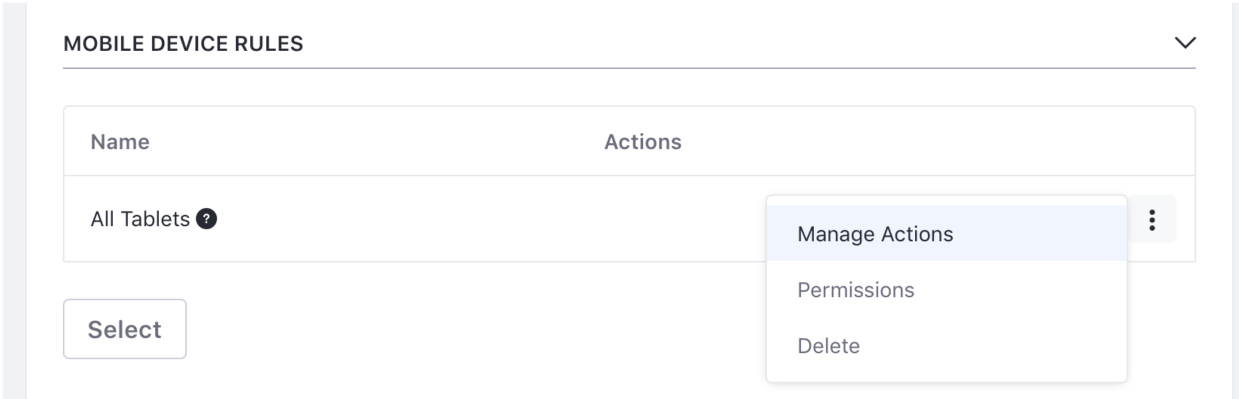

Figure 11.4: Getting to the Manage Actions page.

By default, there are four kinds of actions that can be configured for mobile families:

Layout Template Modification: Changes the way portlets are arranged on pages delivered to mobile devices. For example, you could have pages with more complex layouts automatically switch to a simpler template if it detects a mobile device—even if the resolution is theoretically high enough to support the standard layout.

Theme Modification: Selects a specific theme for different mobile device families. You'd have to have a mobile version of your Site's theme that is automatically applied when a device hits your page.

URL Redirect: Sends mobile users to any URL. This can be used to direct mobile users to a mobile app download or a mobile version of the page.

Site Redirect: Sends mobile users to a different Site on your portal. In some cases, mobile content could be created on a mirror of your Site.

Tip: 7.0 was designed from the ground up to be responsive and adapt to any device that might be accessing it. Before creating new themes or forcing a layout template change, you should test how the Site behaves using Liferay DXP out of the box. Certain features, like URL Redirects, can be disruptive and frustrating for users if used improperly.

Like mobile device rules, mobile device actions are extensible. Your developers can define custom actions in addition to the four actions provided by default.

To review, if you want to configure an action or actions that take place when mobile device requests are received, take the following steps:

1. Create a mobile device family to represent the group of devices for which to define an action or actions.

2. Define one or more rules for your family that describe the group of devices represented by your family.

3. Apply your family to an entire page set of a Site (all the public pages of a Site or all the private pages) or to a single page.

4. Define one or more actions for your family that describe how requests should be handled.

Mobile Device Rules Example

Now you'll look at an example of using mobile device rules. Suppose you want to create a rule so that when a Site is accessed by an Android or iOS tablet, a different layout is used. To set this up, you must follow the same four steps described above.

First create the Mobile Device Family:

1. Navigate to the *Mobile Device Families* page of *Site Administration*.
2. Click *Add Device Family* ( ).
3. Enter *Android and iOS Tablets* for the *Name*.
4. Click *Save*.

Next create a rule for the family:

1. From the *Mobile Device Families* page, click on *Android and iOS Tablets*.
2. Click *Add Classification Rule* ( ).
3. Name the rule *Rule 1*.
4. Under *Operating System* select *Android* and *iPhone OS* (you can hold or to select multiple items).
5. Under *Device Type* select *Tablets*,
6. Click *Save*.

As with the previous example, you only need one rule to describe your device family. Now you must apply the rule to some pages.

1. Go to *Build* → *Pages* in Site Administration.
2. Click on the *Configuration* icon for the *Public Pages*
3. Go to the *Advanced* tab.
4. Under *Mobile Device Rules*, select the *Android and iOS Tablets* device family.

Now you must define an action for your Android and iOS Tablets rule group to use a different layout.

1. Click *Actions* → *Manage Actions* for the *Android and iOS Tablets* rule.
2. Click *Add Action*.
3. Enter the name *Layout Template Modification*, and select the *Layout Template Modification* action type.
4. Select the *1 Column* layout template.
5. Click *Save*.

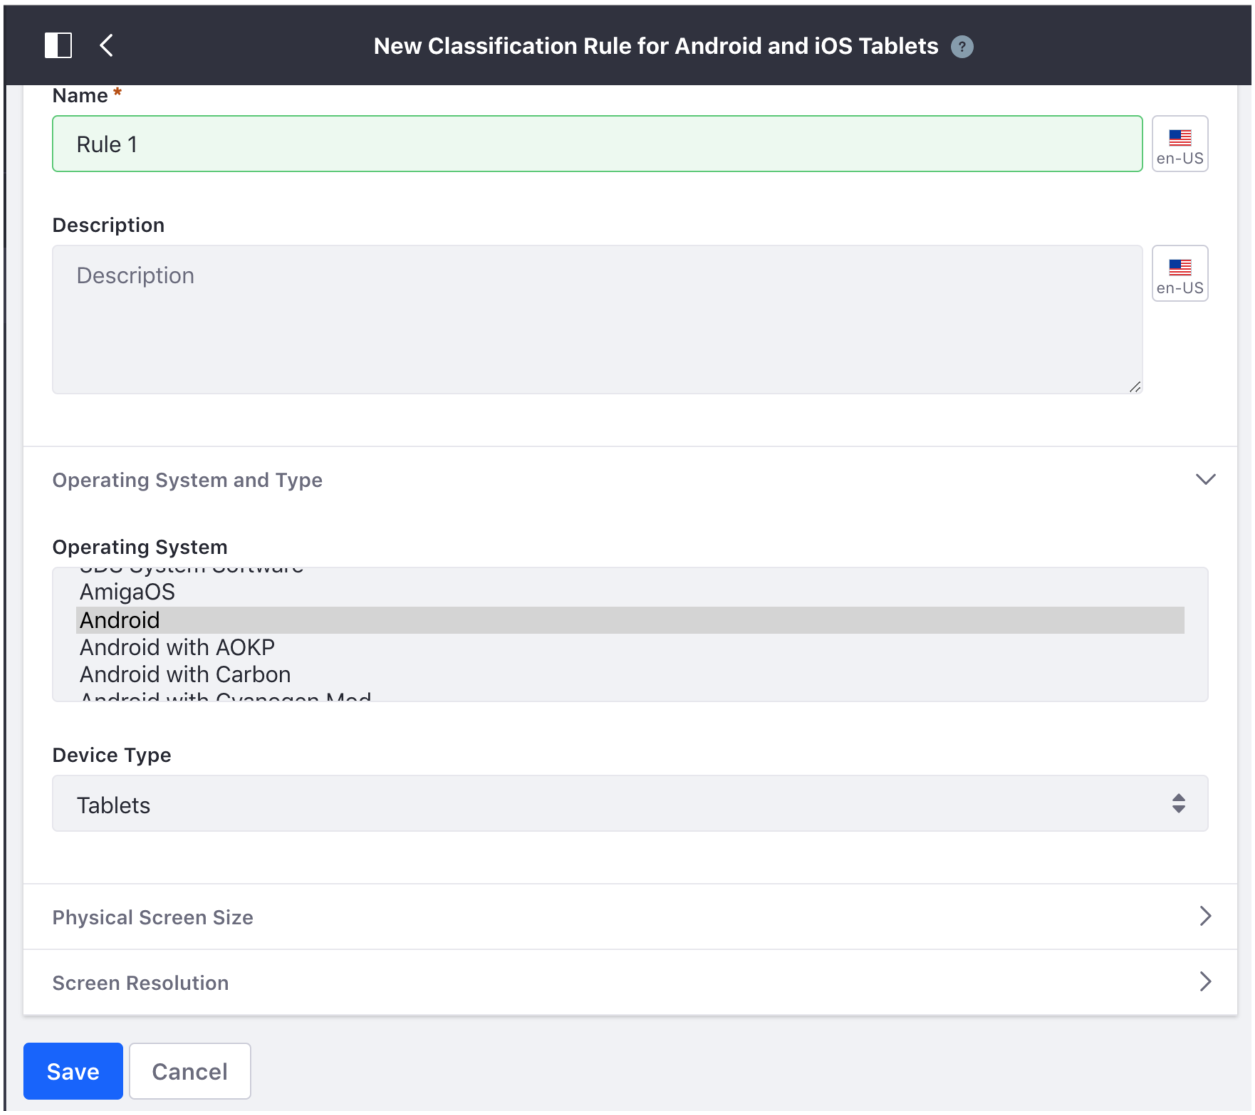

Figure 11.5: Create the Classification rule.

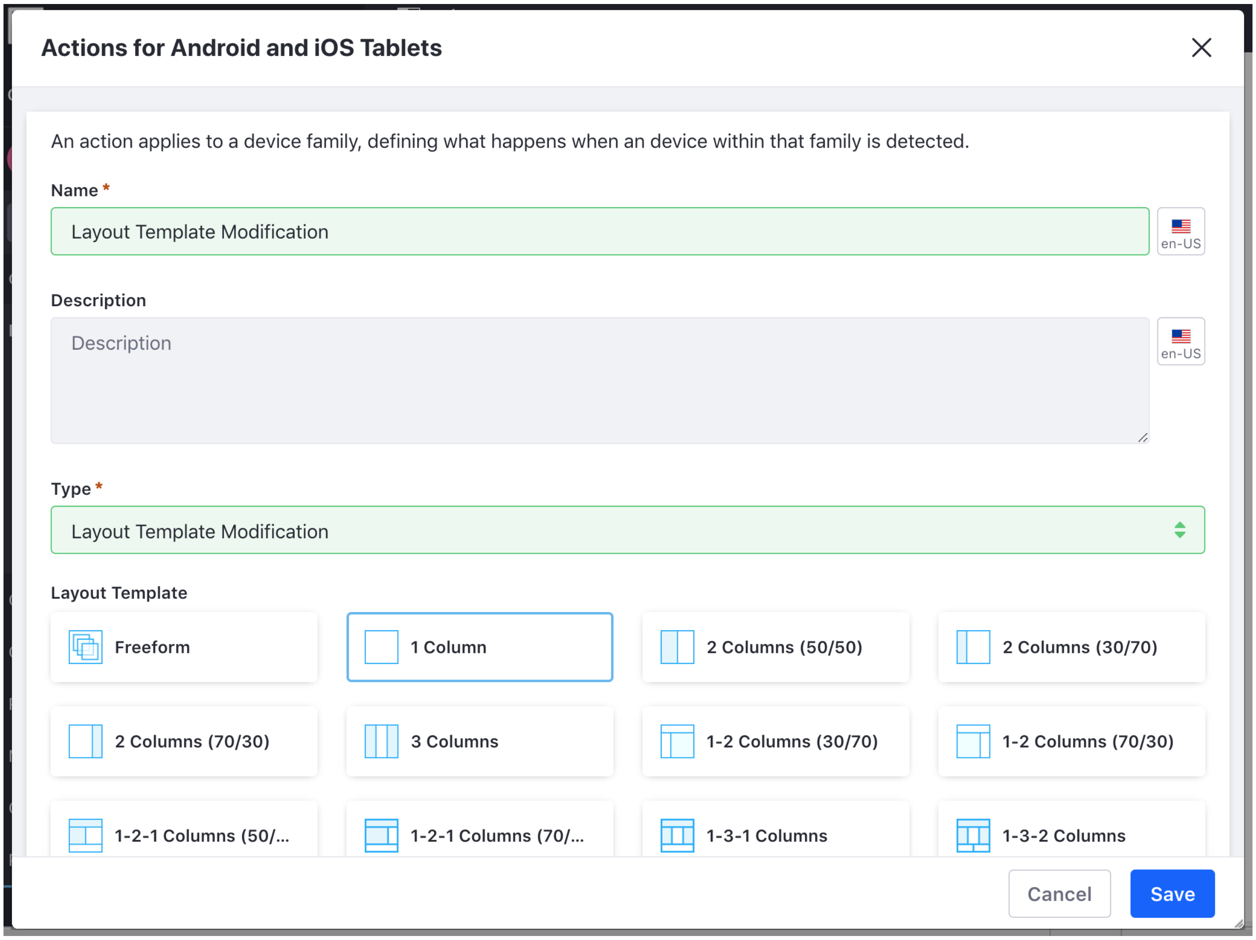

Figure 11.6: Create the Actions for Android and iOS Tablets.

Now the Liferay Site's pages are presented to Android and iOS tablet users with the 1 Column layout template.

Mobile Device Rules are a powerful way to manage the way pages and content appear on the various devices that access your Site. But remember to consider the power of modern devices and the experience of your users, and use this great power responsibly—to help users have a great experience on your website and to not interrupt or negatively impact that experience on whatever device they're using.

Chapter 12

Managing Site Navigation

Liferay provides powerful tools for creating and organizing pages. You can have anything from a simple, flat Site navigation to a complex hierarchy with tree of sub-pages nested down many levels.

7.0 lets you create Navigation Menus separate from your page hierarchy. Now you have the freedom to leave one-off marketing landing pages out of the navigation or create multiple navigation menus: a main menu, secondary menus, footer menus, and custom menus for anything that you can dream up.

Menus can differ by page: landing pages can show a simple list of frequently visited pages, and the rest can appear in secondary navigation. You can also create specific menus for different landing pages to direct users to content that is relevant to them.

Go to *Navigation → Pages* to view the existing pages or create new pages. The Site hierarchy as displayed on *Pages* is the main reference for the organization of pages on that Site. While Navigation Menus can customize their organization and what appears and what doesn't appear, this menu is always the primary reference for the pages on your Site.

12.1 Page Hierarchy

Using the Page Hierarchy, you create public and private pages and organize those pages in whatever order or structure that you see fit.

Creating a Page

New pages are created on the *Build → Pages* page in Site Administration. Pages can be created as *Public Pages* which anyone can view, or *Private Pages* which can only be viewed by Site Members. To create a new page,

1. Go to *Site Administration* for the Site you want to work on, then *Build → Pages*.

2. Click *Add* Add Page and select *Public Page*.

3. Select *Widget Page*.

4. Set the *Name* as *About Us*.

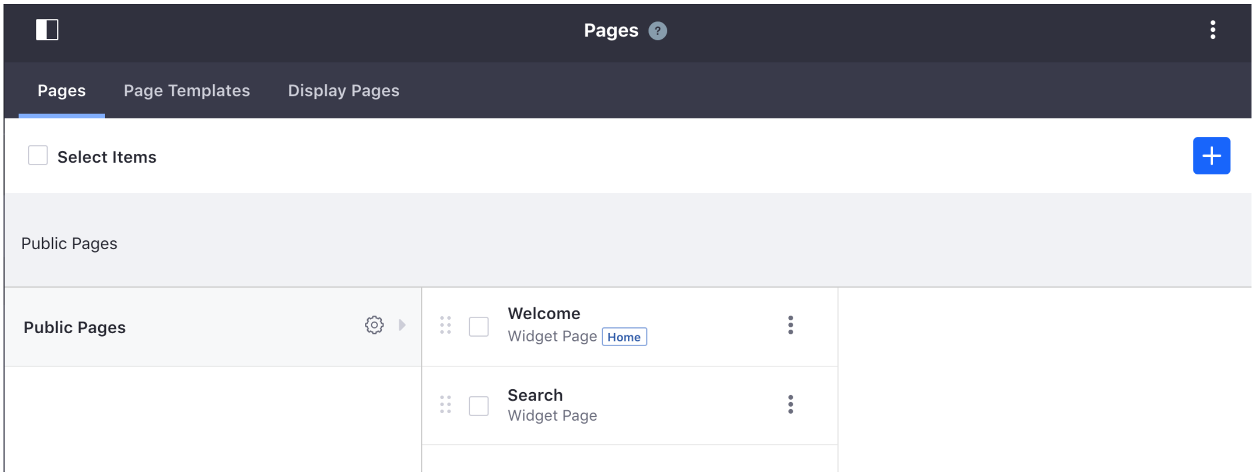

Figure 12.1: In the default site, initially only the *Welcome* and the hidden *Search* pages exist in the Public Pages Hierarchy.

5. Click *Save*.
6. On the next page, note the *Show in Public Page Hierarchy Menu* selection, leave the defaults and click *Save*.

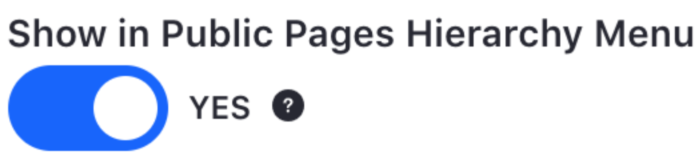

Figure 12.2: When you create a page, by default it is added to the site hierarchy.

Now that the page is created, it appears in the hierarchy, and you can move or organize its position there.

Organizing Pages

Drag and drop pages to reorder their position in the page hierarchy (and subsequently the default navigation that users see), and to nest them as subpages. The page at the top of the list is the *Home* page that users see automatically when visiting your Site.

1. Click *Go to Site* to view the organization of the navigation menu. You can see the order of the pages matches the order of the pages from Site Administration.
2. Drag the *About Us* page above the *Welcome* page in the list. It automatically becomes the *Home* page.
3. Click *Go to Site* to see how this affects your menu.

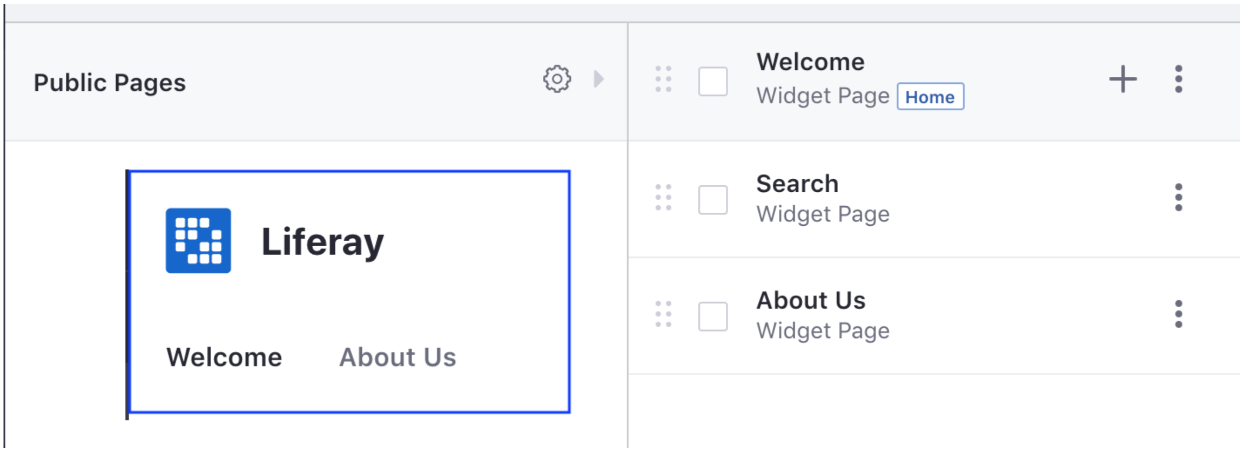

Figure 12.3: You can see the order of pages in Site Administration vs. how they appear on the site.

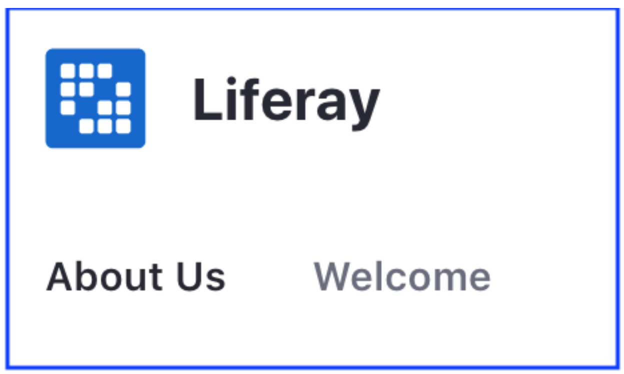

Figure 12.4: *About Us* is now the home page, and *Welcome* is second in the nav.

4. Drag *About Us* on top of *Welcome* to nest it.

5. Click *Go to Site* one more time to see how nested pages appear.

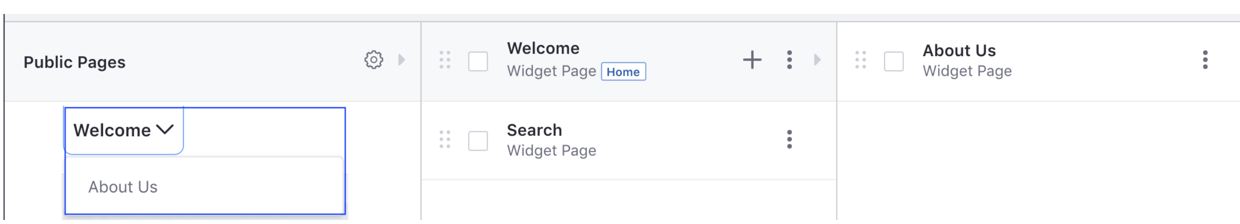

Figure 12.5: *About Us* is now nested under *Welcome* and appear when you mouse-over *Welcome*.

As you've just demonstrated, organizing pages in the default menu is simple, but very powerful.

Public and Private Pages

As noted above, Private Pages work just like Public Pages, except they can be viewed only by registered members of a Site. In the default configuration, Public Pages are at the URL `[web-address]`/**web**/`[site-name]` while Private Pages are at `[web-address]`/**group**/`[site-name]`. Other than the membership distinctions, Public and Private Pages share the same behavior.

Page Options

While managing the default menu, you can also access page options. Clicking on the *Options* icon Option accesses several configuration tools:

- **View** goes to the selected page on the Site.
- **Configure** goes to page configuration.
- **Copy Page** creates a new page in the current Site that duplicates the selected page.
- **Permissions** opens the Permissions dialog.
- **Orphan Widgets** clears data related to widgets that have been removed from the page.
- **Delete** deletes the page and all its data.

Creating and managing pages using the page hierarchy is simple but very powerful. If you need more navigation options, however, Navigation Menus provide more flexibility.

12.2 Creating and Managing Navigation Menus

To better understand Navigation Menus, it's time to create a new menu.

Creating a Navigation Menu

1. Go to *Site Administration* → *Build* → *Navigation Menus*.
2. Click the + button to add a new menu.
3. Give your menu a name and click *Submit*.

On the next page appears a number of elements that you can add to a menu.
Page: Select an existing page from the current Site to add to the navigation menu.
Submenu: Create a second level of menu navigation.
URL: Create a link to any page anywhere by providing a URL. The link appears just like any other option in your menu.
Click on *Page* and you see a view of all the current pages on the current Site. Select a page and click *Add* to add that page to the menu.

Note: When you click on a page, you select that page. Multiple pages can be selected by clicking on each page one at a time. To deselect a page, click on the page again.

Now you see the menu management screen. From here, you can drag and drop menu elements to rearrange or nest them. You can also manage options for this menu by clicking the gear icon in the top right. Let's add another item to the menu.

1. Click the + icon.
2. Select *Submenu* in the menu that pops up.
3. Name your menu *External Links*.
4. Click *Add*.

Click the + button again and select *URL*. You're prompted to enter a page name and URL.

1. Enter a *Name*.
2. Enter the *URL* for an external Site.
3. Click *Add*.

Drag the URL item onto the *External Links* submenu. This nests the URL item in the submenu.

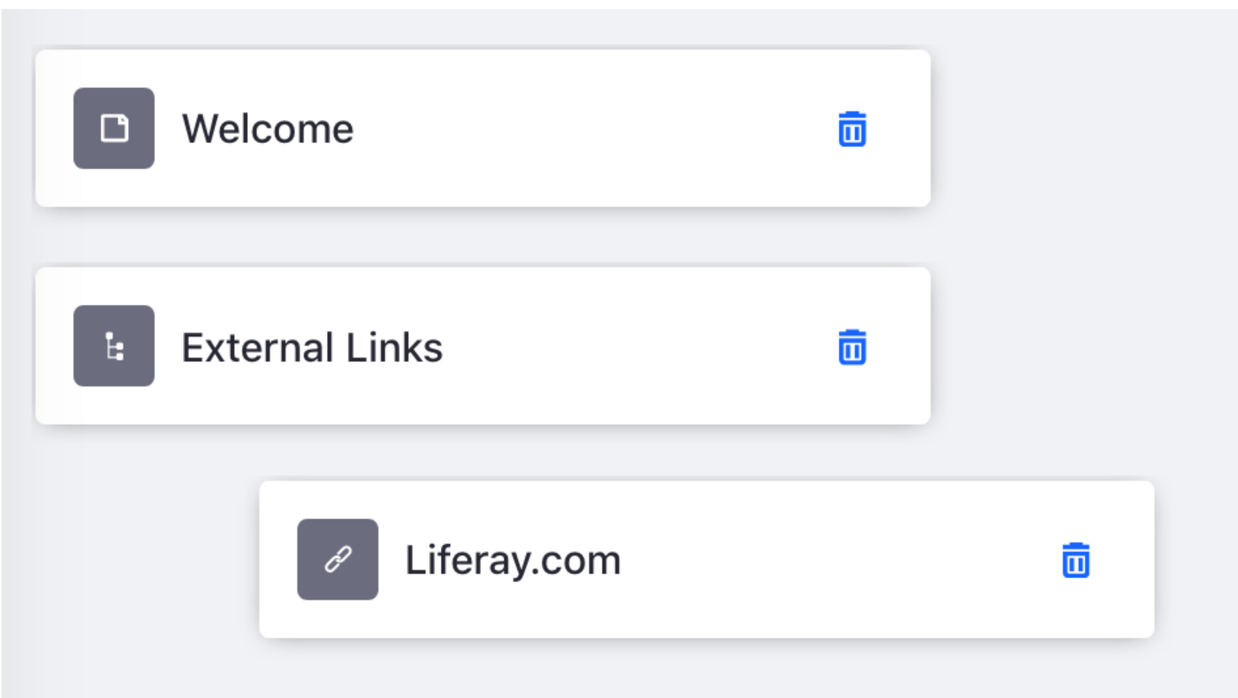

Figure 12.6: Menus can have a standard page, a submenu, and a URL link in the submenu.

Now that you can see how menus work, you can learn the details.

Managing Menus

After you create a menu, more configuration options appear on the main Navigation Menus page.

Title: Your menu's name.

Add New Pages: Determines if new pages added to the main navigation are added to the menu automatically.

Marked As: Can be set as *Primary Navigation*, *Secondary Navigation*, or *Social Navigation*.

Author: The user that created the menu.

Create Date: When the menu was created.

Theme and Fragment developers primarily use the menu types to determine how a menu should be styled. **Primary Navigation** is the main navigation for a page. **Secondary Navigation** is a second level of navigation, possibly a sidebar or a separate menu within a page. **Social Navigation** is for menus that contain links for sharing content on social media or similar tasks.

Modifying Menus

Next click on the options menu at the far left of your new navigation menu:

Edit: Add, remove, or organize menu items.

Rename: Change the name of your menu.

Permissions: Define who can view, update, delete, and manage the permissions for the menu.

Mark As: Change the menu type.

Delete: Deletes the menu.

Next you'll learn about the tools for displaying menus on pages.

12.3 Displaying Navigation Menus

You can display Navigation Menus in different ways on your Site. You may want to configure different display styles for a main menu, sidebar, and footer menu all on one page.

The Navigation Menu Widget

The Navigation Menu widget lets you add navigation wherever you need it. You can place the widget on a page and then select a menu and style for the menu you are displaying.

1. Go to a Widget Page, open the Add menu on the right side of the page and add the *Content Management → Navigation Menu* to the page.

2. Open the *Configuration* menu.

   From here you can configure three main categories:

   - The Navigation Menu to be displayed

   - The styling of the menu

   - What level of navigation to display

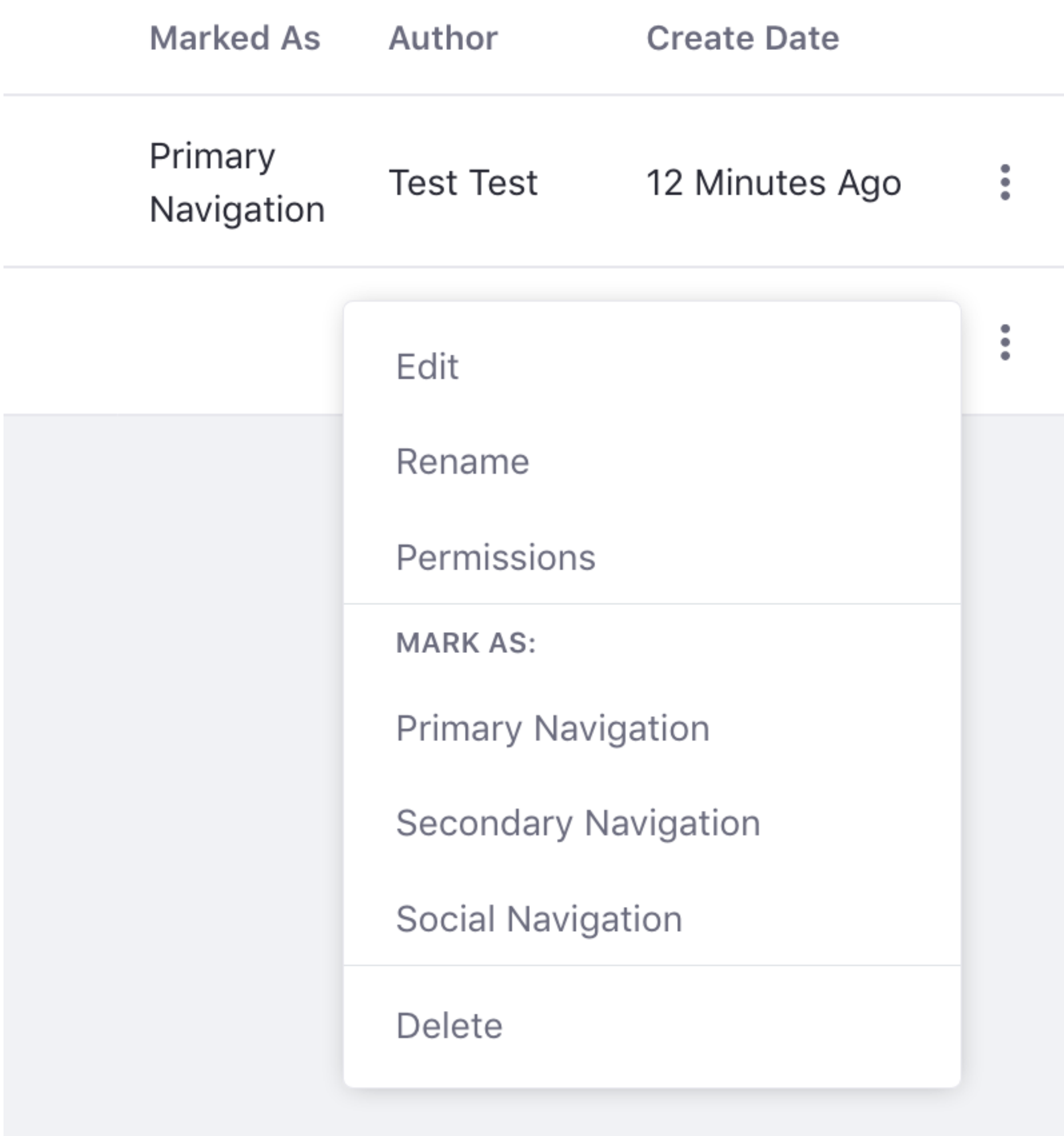

Figure 12.7: Menus with a standard page, a submenu, and a URL link in the submenu are created for different reasons.

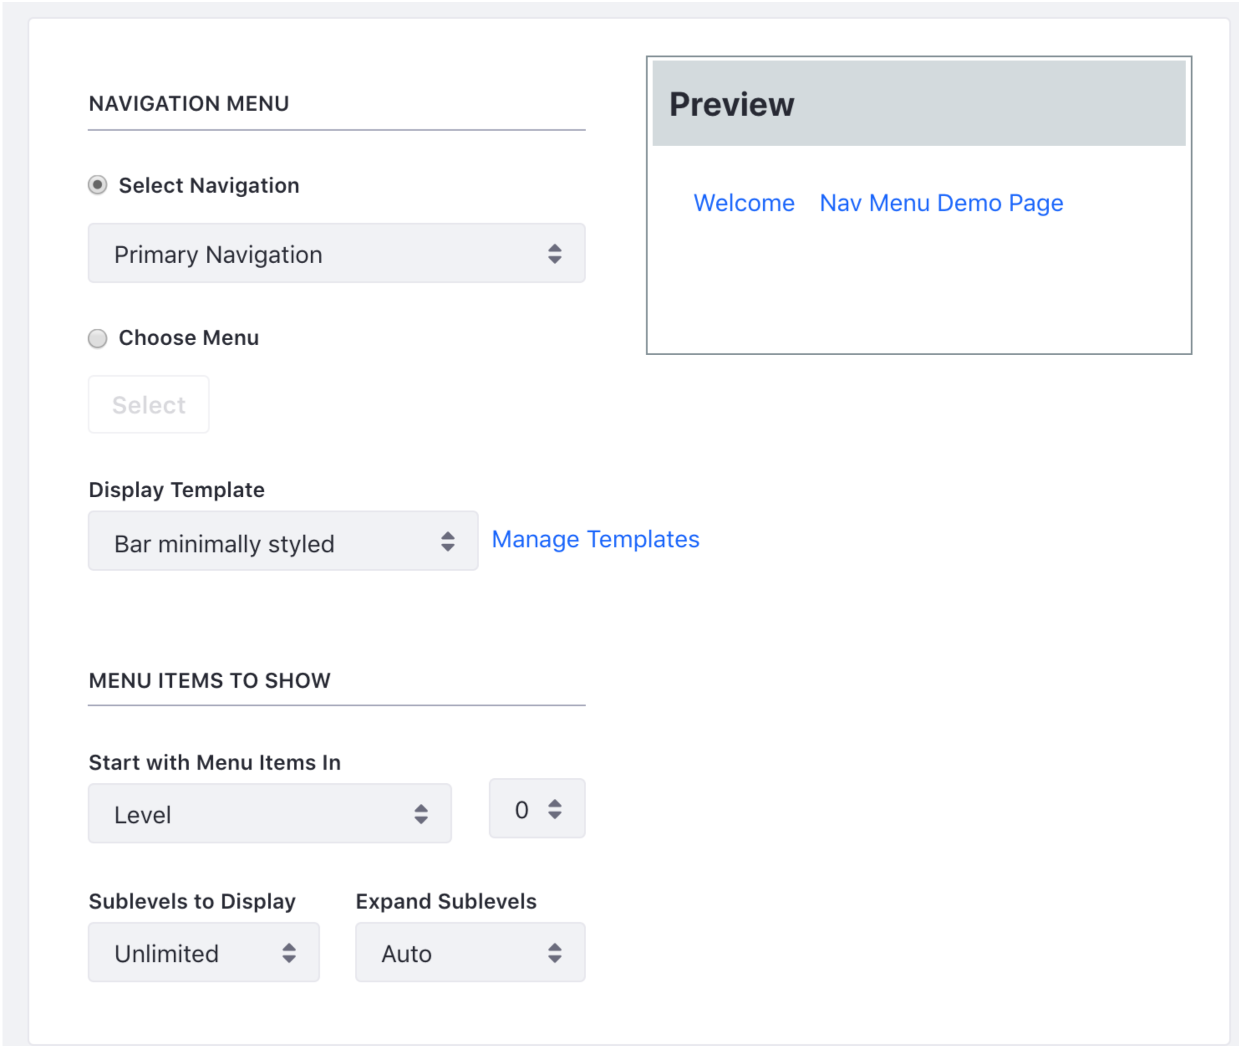

Figure 12.8: Configuring the Navigation Menu Widget.

Choosing a Navigation Menu

The Navigation Menu Widget has two ways to select a menu. You can choose to *Select Navigation* or *Choose Menu*.

Select Navigation: Select from the three main menu types: *Primary Navigation*, *Secondary Navigation*, and *Social Navigation*.

Choose Menu: Choose any menu that was created for that Site.

Once you select a menu, you must choose how to display it.

Display Template

The *Display Template* option lets you select an Application Display Template for Navigation Menus. There ten included by default:

List Menu: Displays all the items in a vertical list.

Pills Horizontal: Displays the items horizontally and uses a button style for highlighting.

Pills Justified: Like Pills Horizontal, but pads the items to fill out the horizontal space.

Pills Stacked: A vertical version of the pills style.

Tabs: Displays the items like navigation tabs.

Tabs Justified: Navigation tabs that fill horizontal space.

Bar Minimally Styled: A lightweight version of the default display that you see in the embedded menu on your page.

Bar Minimally Justified Styled: Like Bar Minimally Styled with horizontal padding.

Bar Default Styled: The default embedded menu.

Split Button Dropdowns: Displays each item as a button with a dropdown for multiple navigation levels.

You can also add your own custom templates.

Menu Items to Show

Menu Items to Show configures which pages at what level from the menu are displayed in the widget. You can choose the starting level, how many levels deep to display, and how to display sub-levels.

Start with Menu Items In: Choose to start at a specific level of the navigation or a level relative to the current level (above or below).

Sublevels to Display: Select the number of levels to display in the navigation, from **1** down to **Infinite**.

Expand Sublevels: Choose if hovering your mouse over the navigation reveals navigation levels one at a time automatically or reveal all the levels at once.

Now you can see how there are a variety of customizations and configurations available for navigation menus that you can implement for your Site.

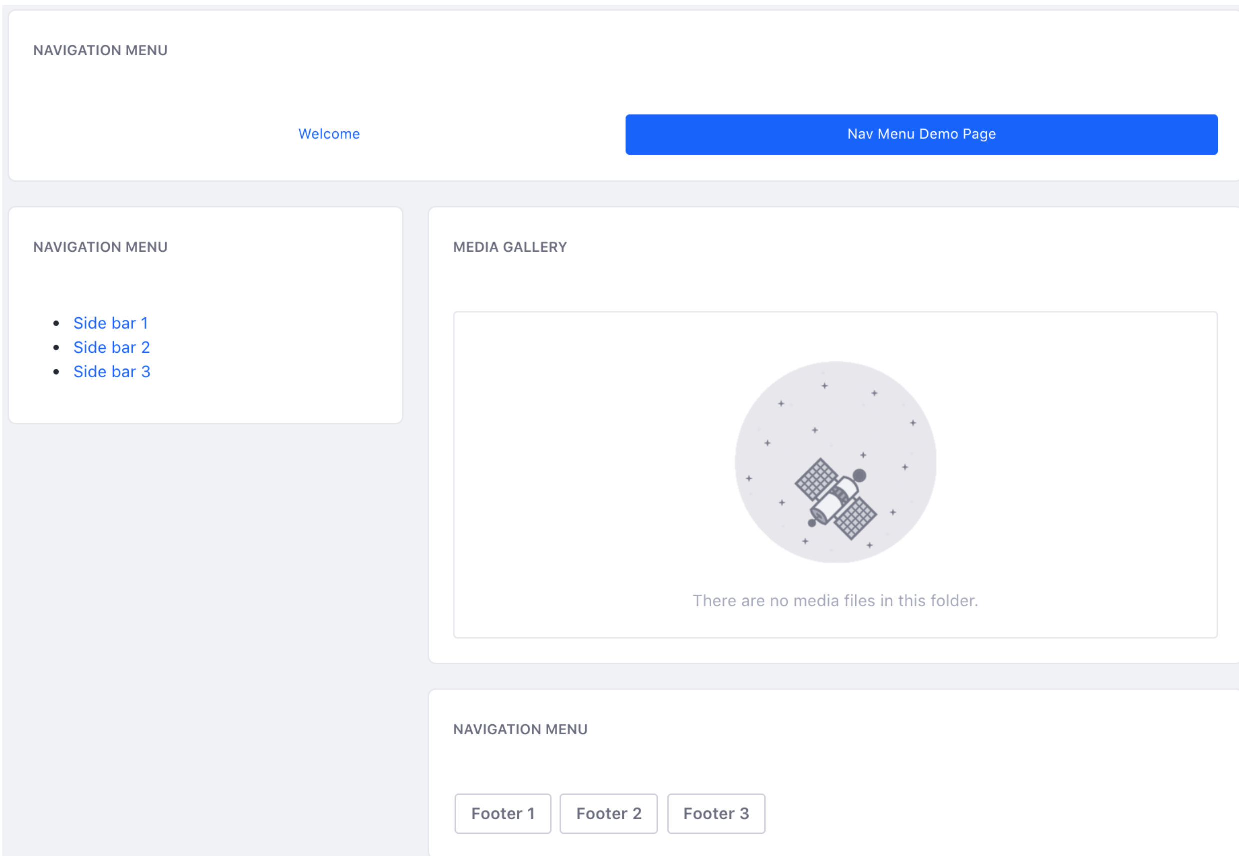

Figure 12.9: Navigation menus give you many ways to help users navigate your Site.

CHAPTER 13

BUILDING SITES FROM TEMPLATES

Site Templates create a single Site structure that can be used for any new Site. They are created and administered from the Control Panel. In addition to creating multiple Sites with the same design, you can also use them to manage changes across multiple Sites with propagation of changes. Site templates can contain multiple pages, each with its own theme, layout template, applications, and app configurations.

Site templates can also contain content just like actual Sites. This allows administrators to use Site Templates to create new Sites that are each created with the same default pages, applications, and content. After they've been created, these Sites and their pages can be modified by Site administrators. Using Site templates can save Site administrators a lot of work even if each Site that was created from a given Site Template ends up being very different.

13.1 Creating a Site Template

Suppose you need to create the following three private Sites for the Lunar Resort's internal use: Engineering, Marketing, and Legal. These should be accessible only to members of these respective departments. You could design each Site separately, but you can save yourself time if you create a Site template instead.

1. Go to the Control Panel and click *Sites* → *Site Templates*.
2. Click the *Add* icon (+) and enter the name *Department* for your template.
3. Leave the *Active* and *Allow Site administrators to modify pages associated with this Site template...* boxes checked.
4. Click *Save* to create your Site template.

The *Active* box must be checked for your template to be usable. If your template is still a work in progress, uncheck it to ensure that no one uses it until it's ready. Checking *Allow Site administrators to modify pages associated with this Site template...* allows Site administrators to modify or remove the pages and apps that the template introduces to their Sites—if you want the templates to be completely static, you should uncheck this.

Now it's time to edit your Site template. This example, includes four pages.

This section assumes knowledge of Liferay 7.1 page management. For more information on how to create and manage pages in Liferay 7.1, see the Adding Pages to Sites tutorial.

---

1. Click the *Options* icon ( ) and select *Manage*.

   This brings you to the *Pages* page for the Site Template. You already have a home page. Create three more pages.

2. Create a *Full Page Application* page named *Documents*.

3. Click *Options* ( ) → *Configure* and set the *Full Page Application* to *Documents and Media*.

4. Create a page using the Global Page Template *Wiki* and name it *Wiki*.

5. Create a widget page named *Message Boards*.

6. Click *Go to Site* in the menu to the left to go to the pages you just created.

7. On the *Home* page add the Activities, Announcements, and Calendar apps.

8. On the *Message Boards* page add the Message Boards and Tag Cloud apps.

The changes you made to your Site template above are completed in real time, so there's no *Save* button.

Next, you'll use your Site template to create the Engineering, Marketing and Legal Sites.

1. Go to the Control Panel and click on *Sites* → *Sites*.

2. Click the *Add* icon ( ) → *Department*.

3. Enter *Engineering* for the Site name.

4. Check the *Create default pages as private (available only to members). If unchecked, they will be public (available to anyone)* option since the Engineering Site is intended for internal use only.

5. Click *Save*.

6. In the next section, set the Membership Type to *Private*.

   Recall that private Sites don't appear in the My Sites application so that regular users won't even know that the Engineering Site exists. Also, the only way users can be added to a private Site is via an invitation from a Site administrator.

7. Leave the *Active* selector enabled so that your Site can be used immediately.

8. Check the *Create default pages as private (available only to members). If unchecked, they will be public (available to anyone).* option since the Engineering Site is intended for internal use only.

9. Leave the *Enable propagation of changes from the Site template* box enabled so that the Engineering Site receives updates if the Department Site template is modified.

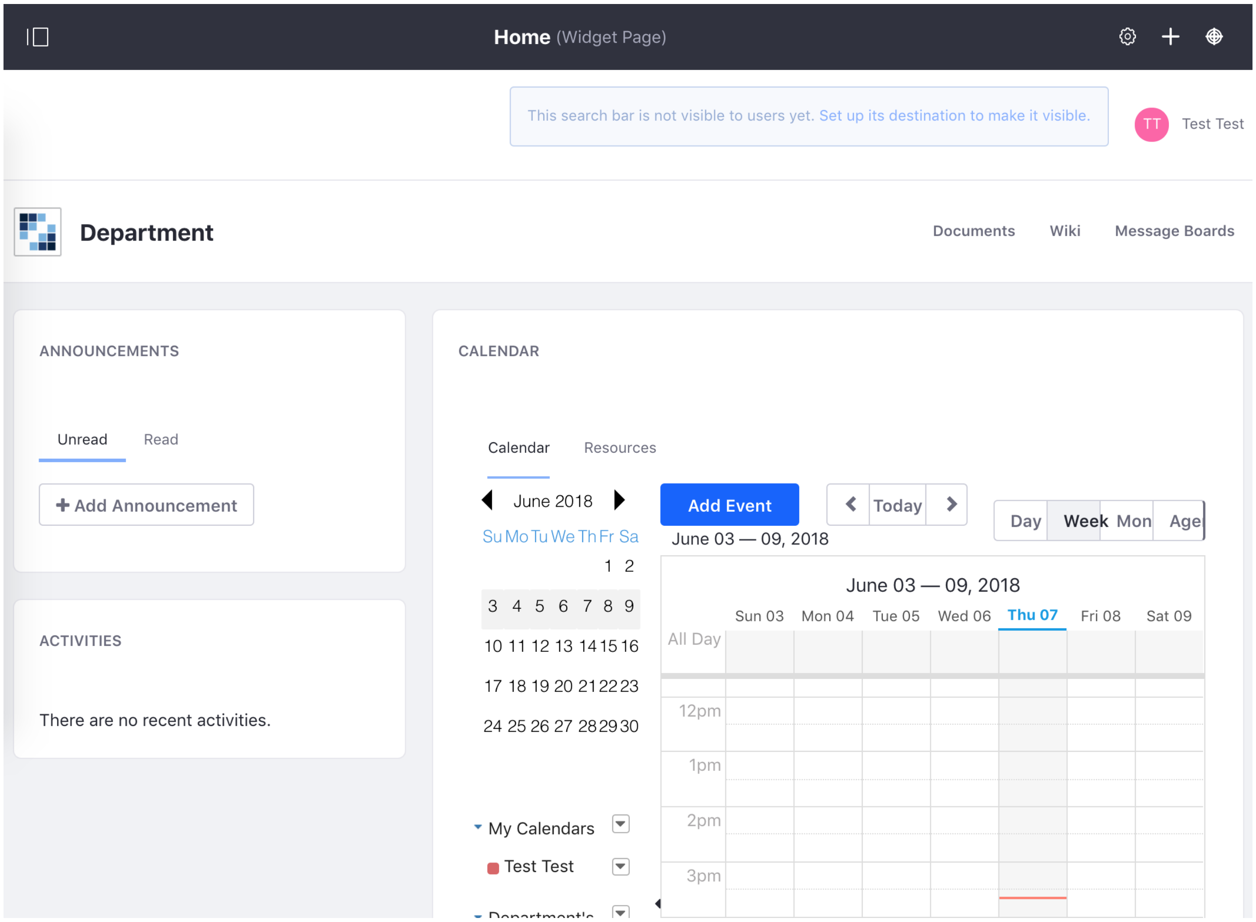

Figure 13.1: You can see the name of the Site template you're currently editing.

10. Click *Save* to create your Engineering Site.

11. Repeat these steps to create the Marketing and Legal Sites.

The new Sites have all the pages and apps you created in the Site template. To view the pages of the new Sites, click *Sites* → *Sites* in the Control Panel and then click *Actions* → *Go to Private Pages* next to one of your new Sites.

Using Site templates streamlines the Site creation process for administrators, making it easy to create Sites quickly. Now each department of the Lunar Resort has its own Calendar, Documents and Media Library, Wiki, and Message Boards on their Sites. Although the pages and apps of each department's Site are the same, each Site will quickly be filled with department-specific information as users add and share content within the Sites. Also, Site administrators can add new pages, apps, and content to their Sites, further differentiating each department's Site from the others.

13.2 Managing Site Templates

To get started, click on *Site Templates* in the Sites section of the Control Panel. Here, you can add, manage, or delete Site templates. You can also configure the permissions of Site templates. As long as a Site is linked to the Site template it was created from, changes to the Site template's pages, apps, and app configurations are propagated to the Site. Changes to a Site template's content, however, are not propagated to existing Sites that are linked to the Site template. You'll learn about the propagation of changes between Site templates and Sites in more detail in the section on Site template use cases below.

To manage a Site Template's pages,

1. Click on *Site Templates* in the Control Panel.

2. Select the *Actions* icon (⋮) and then *Manage* for an existing template.

If you open the main Menu on the left side of your screen (if necessary), the Site Template is selected in the Site Administration dropdown menu. You're provided similar options as a regular Site, including *Build*, *Content*, *Configuration*, and *Publishing*. By default, the Manage Interface opens *Build* → *Pages*. From here, you can add or remove pages from a Site Template or select themes and layout templates to apply to the Site Template. You can also configure each page to have any theme, any layout template, and any number of applications, just like a page of a regular Site. As with Site pages, you can organize a Site Template's pages into hierarchies. When you create a Site using a Site template, the configuration of pages and apps is copied from the template to the Site. By default, all changes made to the Site template are automatically copied to Sites based on that template.

Tip: If you want to publish a piece of web content to many Sites and ensure modifications are applied to all, don't use Site template content for that purpose. Instead, place the content in the global scope and then reference it from a *Web Content Display* application in each Site.

The Content section offers separate repositories for content related apps based on your Site Template. For instance, by clicking *Polls* from the Content section, you can create a poll question that is only available for that specific Site template. Assets created within your template's Content section can only be accessed by Sites using the template.

The Configuration section includes Application Display Templates and Mobile Device configuration options for your Site Template. Also, nested in the Configuration section is the *Site Template Settings*. This edits the template's name and description while also offering boolean options for activating your Site template and allowing Site administrators to modify pages associated with your template.

The following figure displays the form shown when editing the *Department* template's settings: By default, the following Site templates are provided:

- **Intranet Site:** Provides a preconfigured Site for an intranet. The Home page displays the activities of the members of the Site, search, a language selector, and a list of the recent content created in the intranet. It also provides two additional pages for Documents and Media and external News obtained through public feeds.

- **Community Site:** Provides a preconfigured Site for building online communities. The Home page of a *community Site* provides message boards, search, a display of a poll and statistics of the activity of community members. The Site will also be created with a page for a wiki.

Now that you know the basics for creating and managing your Site templates, you can learn about propagating changes next.

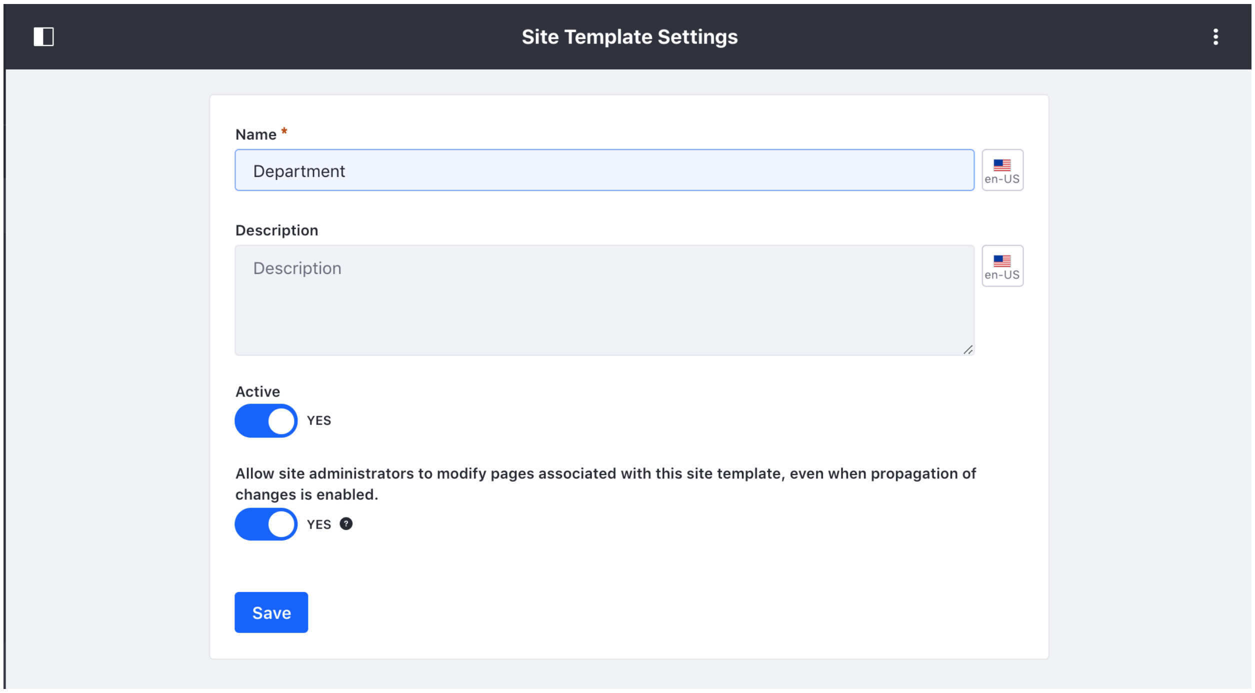

Figure 13.2: Site templates have several configurable options including the option to allow Site administrators to modify pages associated with the Site template.

13.3 Propagating Changes from Site Templates to Sites

Site Template administrators can add, update, or delete Site Template pages. Changes made to a Site Template can be propagated to Sites whose page sets are linked to the Site Template. When you create a Site based on a Site Template with the *Enable propagation of changes from the Site template* box checked this link is created. To configure propagation of changes:

1. Select the Site from the Sites dropdown in the Menu by selecting the *Site Selector* button ().

2. Navigate to the *Configuration* → *Site Settings* page and uncheck or recheck the *Enable propagation of changes from the Site template* checkbox.

In this section, you'll learn about the propagation of changes from Site templates to Sites and discuss the options available to Site administrators and Site template administrators.

Site Template Page Behavior

If a Site's page set has been created from a Site template and the propagation of changes from the Site template is enabled, Site administrators can add new pages but cannot remove or reorder the pages imported from the Site Template. If a Site has both pages imported from a Site template and custom Site pages, the Site Template pages always appear first in the Site page hierarchy; custom pages added by Site administrators appear after the Site template pages. Only Site template administrators can remove, reorder, or add Site template pages. Site administrators can add or remove custom Site pages. They can also reorder custom Site pages as long as they're all positioned after the Site template pages. Site template administrators cannot add, remove, or reorder custom Site pages.

Note: Pages containing a fragment (e.g., Content Pages) cannot propagate changes after a Site is *first* created based on a site template.

If a Site administrator changes a page that was imported from a Site Template and refreshes the page, the following Information icon (i) appears in the Control Menu with the following message:

```
This page has been changed since the last update from the Site template. No
further updates from the Site template will be applied.
```

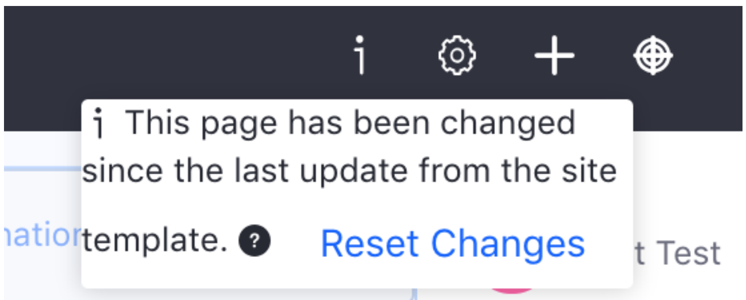

Figure 13.3: You can click the Information icon to view important information about your Site template.

Merging and Resetting Changes

If the Site administrator clicks the *Reset Changes* button, changes are propagated from the Site template page to the corresponding Site page that was imported from the Site template. Clicking the *Reset Changes* button makes two kinds of updates to a page. First, changes made by Site administrators to the Site page are undone. Second, changes made by Site template administrators to the Site template page are applied to the Site page. Note: clicking the *Reset Changes* button only resets one page. If multiple Site pages have been modified and you'd like to re-apply the Site template pages to them, you'll need to click the *Reset Changes* button for each page.

Site template administrators can set preferences for apps on Site template pages. When a Liferay administrator creates a Site from a Site template, the app preferences are copied from the Site template's apps, overriding any default app preferences. When merging Site template and Site changes (e.g., when resetting), app preferences are copied from Site template apps to Site apps. Only global app preferences or local app preferences which don't refer to IDs are overwritten.

In some cases, merging Site template and Site changes fails. For example, if pages from a Site template cannot be propagated because their friendly URLs are in conflict, Liferay DXP could try to continuously merge the Site changes. Instead of entering into an infinite loop of merge fails, Liferay DXP stops the merge after several unsuccessful attempts. Liferay DXP, however, doesn't stop there: your merge is temporarily paused, you're given an indication of the current merge fail, and then you have the opportunity to fix your merge conflicts. After you've squared away your conflict, navigate to your Site's *Site Administration* → *Configuration* → *Site Settings* and click the *Reset and Propagate* button.

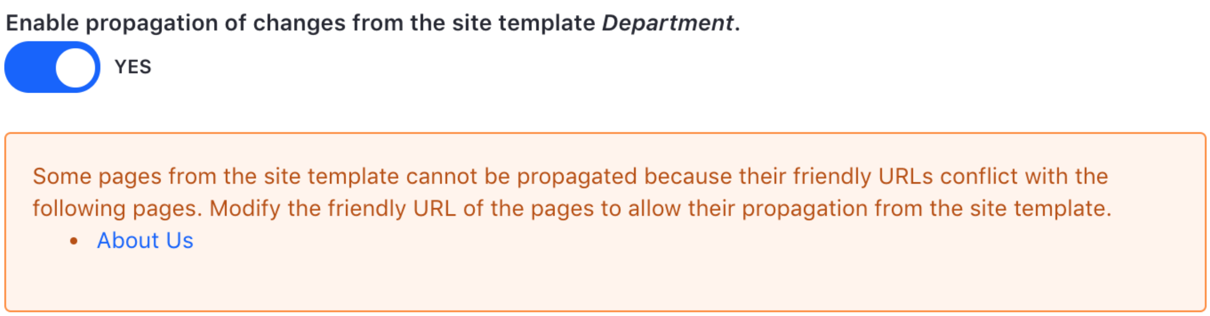

Figure 13.4: This type of warning is given when there are friendly URL conflicts with Site template pages.

The *Reset and Propagate* button resets the merge fail count and attempts to propagate your Site changes again. This process gives you the opportunity to detect and fix a merge fail when problems arise. This helpful process can also be done with page template merges, which follows similar steps.

Site administrators can also add data to Site template applications. For example, Site template administrators can add the Wiki app to a Site template page and use the Wiki to create lots of articles. When a Liferay administrator creates a Site from a Site template, data is copied from the Site template's apps to the Site's apps. The preferences of the Site's apps are updated with the IDs of the copied data. For example, if a Site is created from a Site template that has a Wiki app with lots of wiki articles, the wiki articles are copied from the Site template's scope to the Site's scope and the Site's Wiki app is updated with the IDs of the copied wiki articles.

Important: App data, fragment-based pages, related resources, and permissions on resources are only copied from a Site template to a Site when that Site is *first* created based on the template. No changes made to these entities are propagated to the Site after the Site is created. Neither are such changes propagated to a Site by the *Reset* or *Reset and Propagate* features.

For example, consider a Site template administrator who includes a Message Boards app as part of a Site template. They even create Message Board categories and configures permissions over the actions of the categories. The first time a Site is created based on the Site template, the categories (app data) and related permissions are copied to the Site. If the Site template administrator adds, removes, or deletes some categories, however, such changes *aren't* propagated to the Site.

Now that you've learned how Site templates work, you'll learn how to share Site templates.

13.4 Sharing Site Templates

If you want to export a Site that uses Site or Page Templates to a different environment (through a LAR file or remote publication), the templates must be exported and imported manually in advance or the import fails.

To export a Site using a Site Template, use the following process:

1. Go to the *Control Panel* → *Sites* → *Site Templates* menu.

2. Click the *Actions* icon ( ) and then *Manage* for the Site template you want to export.

3. Open the Site Template's management section and click on *Publishing* → *Export*.

4. Click the *Add* icon () to create a new Custom Export.

5. Select the content and pages you want to export and click *Export*.

6. Click on the *Download* icon for the template that you exported.

7. In your target environment, go to *Control Panel* → *Sites* → *Site Templates* and create a new Site template.

8. Click *Actions* → *Import* for that Site template and upload the LAR file containing your Site template's content.

Now the template can be used normally in the new environment. For more information on exporting/importing content, visit the Importing/Exporting Pages and Content article.

CHAPTER 14

CONFIGURING SITES

Just like there's more than one way to cook and an egg — or eat an egg for that matter — there's more than one way to build a site. Liferay is created with that in mind, and beyond just creating content and pages, Liferay has a wealth of configuration options and tools available to help you create the site that meets your needs and the needs of your users.

In this section, you'll explore all of Liferay's options for configuring sites, and gain a deeper understanding of why you might want to use various options and configurations.

14.1 General Settings

General settings range from core configuration, like your Site's Membership Type, to finer details like Documents and Media indexing options.

Details

Details provides the same menu you filled out when first creating your Site. This allows an administrator to change the description and membership type of a Site.

Membership Options

The membership type can be set as open, restricted, or private based on the privacy needs of the Site. Users can join and leave an open Site at will. To join a restricted Site, users must be added by the Site administrator, but they can request membership through the Sites section of the Control Panel. A private Site works like a restricted Site but is hidden from users who aren't members.

Site Hierarchies

Sites can be organized into hierarchies. At the bottom of the Details sub-section is the Parent Site section. When you select the parent Site for the Site you're currently on, a checkbox appears for limiting membership to members of the parent Site.

Pages

Under Pages you can view your Site's Public or Private Pates, if any exist. If they don't exist, a *Site Templates* selector appears for creating pages with a Site Template.

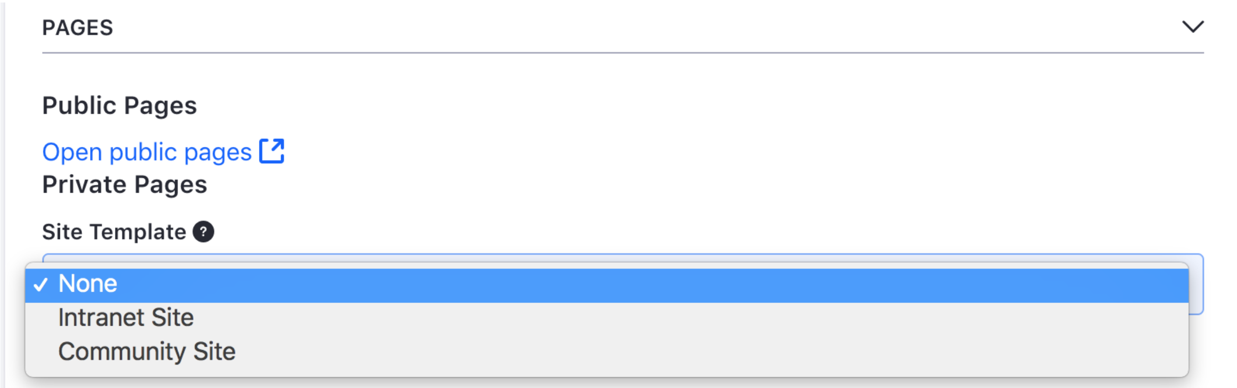

Figure 14.1: Selecting a Site Template.

Categorization

Categorization helps administrators organize the Site and allows for users to easily find your Site and its content through search and navigation. For more information on using tags and categories, visit the Organizing Content with Tags and Categories section.

Site URL

The *Friendly URL* option lets you set your Site's URL paths. Friendly URLs are used for both public and private pages. The public Site base URL is `https://localhost:8080/web`, and the private one is `https://localhost:8080/group`. Each friendly URL must be unique.

For example, setting the friendly URL of your default Site to `/lunar-resort` makes your Site's public home page's URL `https://localhost:8080/web/lunar-resort/home`. The private Site's URL is thus `https://localhost:8080/group/lunar-resort/home`.

Note that if you add a friendly URL for your instance's home page, you should update your instance's Home URL field so that page requests to `http://localhost:8080` redirect properly:

1. Go to *Configuration* → *Instance Settings* in the Control Panel.

2. Under *Navigation*, in the Home URL field enter your home URL (i.e. */web/lunar-resort/home*).

Once you've entered this setting, page requests to `localhost:8080` redirect to the friendly URL of your Liferay instance's new home page.

You can also configure Virtual Hosts, which connects a domain name to a Site, under *Site URL*. You can use this to define a domain name (i.e., `www.lunar-resort.com`) for your Site. This can be a full domain or a subdomain. You can use this to host a number of web sites as separate Sites on one Liferay server.

For instance, if you set this up for the Lunar Resort's development network, users in that Site would access `developers.lunar-resort.com`, provided that the Lunar Resort instance's network administrators created the domain name and pointed it to the Liferay server.

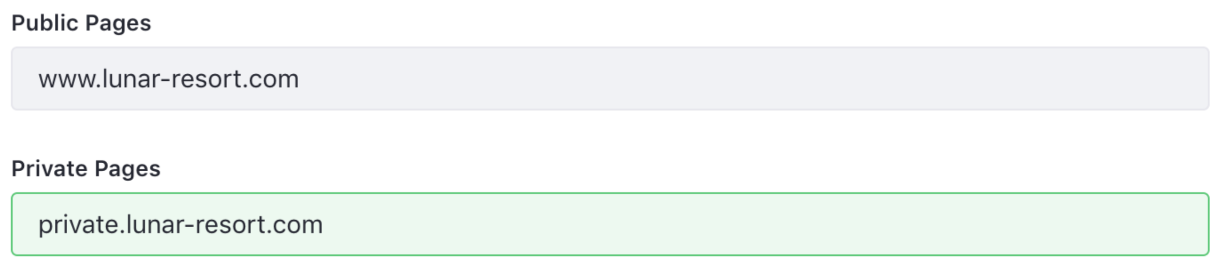

Figure 14.2: When configuring virtual hosts, the public and private pages of a site can be configured to different domains.

1. With your provider, set the DNS name *developers.lunar-resort.com* to point to your Liferay instance's IP address.
2. In the Virtual Host tab for the Developers Site, set the URL to *http://developers.lunar-resort.com*

This helps users quickly access their Site without having to recall an extended URL. The *Site URL* option is listed under the General tab.

Documents and Media

The last option is *Documents and Media*, which lets you enable/disable Directory Indexing. If on, Site administrators can browse your Site's Documents and Media files and folders. For example, a Site administrator of a Site called *Lunar Resort* can browse documents at `http://localhost:8080/documents/lunar-resort` if this option is enabled.

Site Template

If you created your Site using a Site Template, this section appears and displays information about the link between the Site Template and the Site. Specifically, you can see which Site Template was used and whether or not it allows modifications to the pages inherited from it by Site administrators. To learn more about Site Templates and how to create your own, see Building Sites from Templates.

Custom Fields

Custom Fields only appears if you've created them in Control Panel → *Configuration* → *Custom Fields*. For more information on Custom Fields, see Custom Fields.

14.2 Social Settings and Languages

The Social tab provides options for managing the social interactions on your Site. Languages lets you configure language options and change the default language options for the Site.

Ratings

The *Ratings* option lets you select the ratings type to use for applications like Documents and Media, Web Content, Comments, etc. Ratings types include Stars, Likes, and Thumbs.

Mentions

At the bottom of the page is *Mentions*, which lets you enable/disable Mentioning functionality, which is used to *mention* (notify and/or draw attention to) friends and colleagues by entering the "@" character followed by their user names. See the Mentioning Users article for more information.

The next tab is *Languages*.

Languages

The *Languages* tab lets you configure the language options for your Site. You can use the default language or define another supported language as the default for your Site.

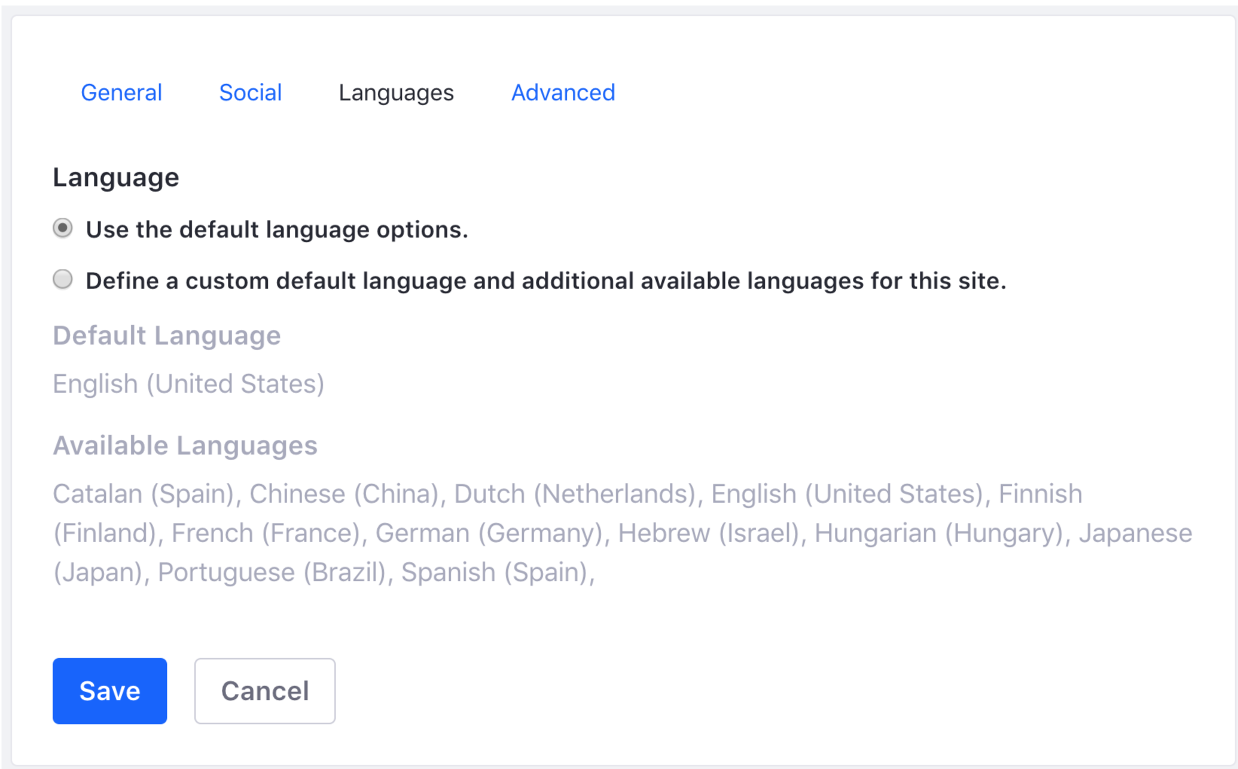

Figure 14.3: In the Languages tab, you can configure the site to use the instance's default language or another supported language.

Now that you know how to configure Sites, you're ready for the advanced settings.

14.3 Advanced Site Settings

Advanced Settings relate to security (like User Roles) or require external configuration (like creating a Google Analytics account) to use.

Default User Associations

Default User Associations configures Site roles and teams that newly assigned Site members have by default. If you'd like to learn more about creating roles and/or teams, visit the Roles and Permissions and Creating Teams for Advanced Site Membership Management.

Analytics

Liferay DXP includes built-in support for Google Analytics for analyzing traffic on your Site. Google Analytics provides a snippet of code which you add to your pages enable tracking. Adding this code to every page on a Site would be tedious, especially if it's a large Site with a lot of user-generated content.

There are two ways to mitigate this problem:

1. A web developer can hard-code the tracking code into a theme, which embeds it on every page.
2. An administrator can enter the tracking code in Site settings.

To use option #2:

1. Go to *Site Settings* → *Advanced*.
2. Expand the *Analytics* section.
3. Enter your Google Analytics ID.
4. Click *Save*.

   All the pages in the Site you selected now have the Google Analytics code and can be tracked.

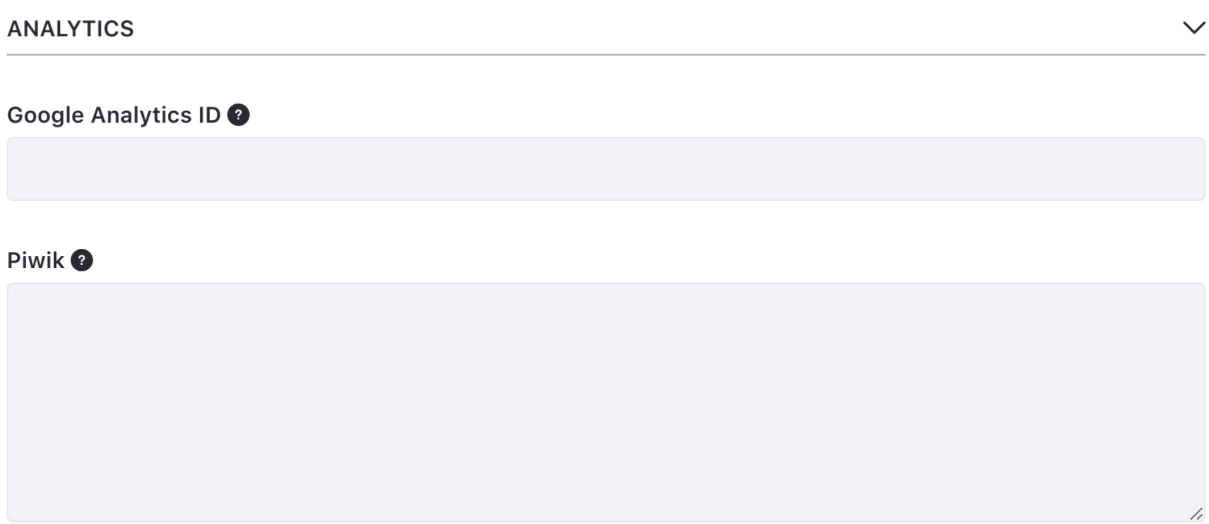

Figure 14.4: To set up Google Analytics: sign up, receive an ID, and then enter it into the Google Analytics ID field.

To enable a different analytics service:

1. Go to *Configuration* in the Control Panel.

2. Go to *Instance Settings → Miscellaneous*.

3. Enter the name of any additional service you want to add in the *Analytics* field provided.

4. Once you have entered the name, go to the *Site Settings → Advanced → Analytics* page for the Site where you wish to add analytics.

5. Copy the JavaScript tracking code provided by your analytics platform into the corresponding field for your service.

Now all pages on the selected Site contain the tracking script and send analytics data to your analytics platform.

Maps

The *Maps* option configures the maps API provider used by your Liferay instance when displaying geolocalized assets. Geolocalized assets can be displayed for documents, web content articles, DDL records, etc. Maps is available under the Advanced tab. You can read more about Geolocation in Geolocating Assets.

Recycle Bin

The *Recycle Bin* option enables or disables the Recycle Bin for your Site. You can also regulate the age (in minutes) for which content is able to be stored in the Recycle Bin until it is permanently deleted. For a full explanation of the Recycle Bin, see Restoring Deleted Assets.

Content Sharing

If you select the *Content Sharing* tab from the Advanced tab, you can configure whether sub-Sites can display content from this Site. Administrators of this Site's sub-Sites can use all structures, templates, categories, application display templates, and more from this parent Site. Even if you initially allowed content sharing between the parent Site and its sub-Sites, you can disable this option and immediately revoke content sharing from all sub-Sites.

14.4 Customizing Personal Sites

By default, newly created users are granted a personal Site.

- Users function as Site administrators of their personal Sites.

- Personal Sites are fully customizable but cannot have more than one member.

- Users can have publicly available content on their Site's Public Pages. This is often used for a user blog.

- Users can also have Private Pages where they can keep personal information or use Documents and Media to have their own private file repositories.

You can disable personal Sites by adding the following properties to your `portal-ext.properties` file:

```
layout.user.public.layouts.enabled=false
layout.user.private.layouts.enabled=false
```

Note: The public and private page sets of personal Sites are handled separately. You can leave one page set enabled while disabling the other.

If you initially had user personal Sites enabled for your instance but then disabled them, existing personal Sites remain on your Liferay instance until the next time users log in, at which point they're removed.

You can allow users to create personal Sites but not have them automatically created for new users. To do this, add the following properties to your `portal-ext.properties` file:

```
layout.user.public.layouts.auto.create=false
layout.user.private.layouts.auto.create=false
```

If the properties `layout.user.public.layouts.enabled`, `layout.user.private.layouts.enabled`, `layout.user.public.layouts.auto.create`, and `layout.user.private.layouts.auto.create` are all set to `true`, which is the default, users have personal Sites and public and private pages are created automatically for new users.

There are a number of portal properties you can use to customize the automatically created pages. You can customize the names of the default pages, the applications that appear on the pages, the themes and layout templates of the default pages, and more. Please refer to the Default User Public Layouts and Default User Private Layouts sections of the `portal.properties` file for details.

Note: By default, users can modify the pages and applications on their personal Sites. Administrators, however, can customize the modifiable portions of personal Sites through Liferay DXP's permissions system by removing permissions from Roles. To disallow all Liferay users from modifying something, remove the relevant permission from the User Role.

14.5 Importing/Exporting Sites and Content

Export/Import lets you backup and restore your Site and app data as a LAR (Liferay Archive). There are two primary places Export/Import is used: Sites and apps. You can learn more about exporting/importing app data in the Exporting/Importing Widget Content section. In this section, you'll learn how to export and import content for Sites.

Backing Up and Restoring Pages and Their Content

In *Site Administration → Publishing*, you can find the *Export* and *Import* option for pages. If you click on *Export*, you see an interface for exporting your public or private pages. The Export feature exports your Site's data as a single LAR file. Similarly, *Import* is a similar interface for importing public or private pages from a LAR file.

When importing data into a Site, you should use a newly created Site to avoid conflicts between the existing data and the data being imported. When exporting Site data, you can specify exactly what data should be included in the LAR:

- Site pages (you can select exactly which ones)
- Page settings
- Theme
- Theme settings
- Logo
- Application configurations
- Application content
- Archived setups
- User preferences

A LAR file can be imported into a Site on another Liferay server. You can take content from a Site in one environment (say, a development or QA environment) and move it all to a Site on another server with LARs. You can use LARs to import data onto production servers, but you should not make this a regular occurrence. If you want to regularly move pages from one server to another, you should use Liferay DXP's staging environment. See Staging Content for Publication.

You can export LARs to use them as a backup. If you ever have to restore your Site, you must only import the latest LAR file. However, please be careful! If there's content that exists both in the LAR and in the Site that's importing the data, there may be a conflict, and data could be corrupted. If you want to restore a Liferay Site using a LAR file, delete the Site entirely, create a new Site with the same name as the old one, and then import the LAR file into the new Site. This way, there's no chance for there to be a data conflict.

Some naming collisions are handled automatically. For example, a collision occurs if the LAR you're importing and the Site both have a page with the same friendly URL. Liferay DXP resolves the collision by adding a number to the end of the friendly URL and incrementing until there's no collision. Similarly, if importing a LAR into a Site causes a category name collision, the imported categories are automatically renamed.

Note: LAR files are version dependent. You can't import a LAR file that was exported from one version of Liferay into a Liferay server that's running a different version of Liferay. Also, note that periodically exporting LARs is *not* a complete backup solution; please refer to the Backing up a Liferay Installation section for information on backing up Liferay.

Page Export Example

Here's how the export process works:

1. Go the *Site Administration* → *Publishing*.

2. Click *Export*.

3. Click *Add* ().

   A *New Custom Export* page loads, so you can choose the pages and content you want to export from your Site.

4. Enter "Lunar Resort Version 1" for the *Title*.

5. Under *Pages*, select public or private pages and the settings you want to export.

6. Under the *Content* category, select *All*.

   Note that if you select one of the *Choose* radio selectors or *Change* links, you're given checkboxes for options to choose. The applications' content can also be selected for export, including the Documents and Media Library, Message Boards, and Web Content assets. You can even export the theme you're using. Finally, you can select whether the permissions for your exported pages and content are included.

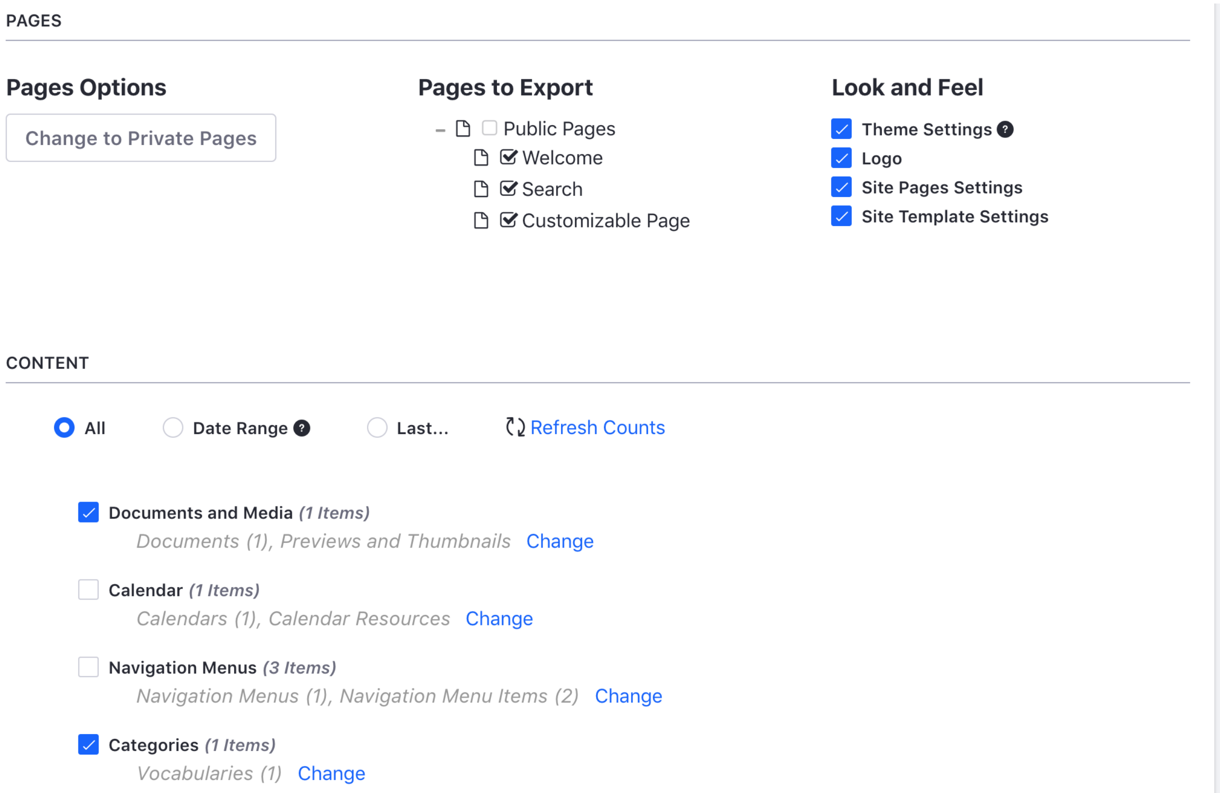

Figure 14.5: You can configure your export options manually by selecting pages, content, and permissions.

7. Click *Export*.

Once you click *Export*, the menu automatically switches to the *Processes* tab, where you see the status of your exported LAR file. You can select the *Download* icon ( ) to download the export to your local machine. Once you have the file, you can copy it to a backup location for safekeeping or import it into another installation of Liferay. If you must rebuild or wish to revert back to this version of your Site, you can import this file by clicking the *Import* button from the Publishing menu, browsing to it, and selecting it. You can also drag a LAR file inside the dotted area, which also executes the import process.

Export Templates

Instead of manually customizing an export process every time you export pages/content, you can use an Export Template. This provides you the convenience of storing export process settings so they can be reused. If you export pages frequently and usually select the same options to export, you can create an export template to export with your standard options.

To create an export template,

1. Select the *Options* icon (⋮) from the top right corner of the screen and select *Export Templates*.
2. Click the *Add* button (+).
3. Assign the template a *Name* and *Description*.
4. Fill out the configuration options for your export process.
5. Click *Save*.

Your template is now available to use from the *Export Templates* menu. To use the template,

1. Click the *Actions* (⋮) next to the template.
2. Select *Export*.

   This automatically fills the fields and options for exporting pages and their content.
3. Give the export a name.
4. Click *Export* and your LAR file is generated.

Chapter 15

Styling Apps and Assets

Application Display Template (ADTs) define custom display templates used to render widgets. For example, you may want to show blog entries horizontally instead of vertically, or list your assets in the asset publisher application in different sizes. In this section you'll learn about the capabilities of ADTs and how to create and configure them.

15.1 Styling Widgets with Application Display Templates

Suppose you're customizing the Lunar Resort Site and want to allow users to use Facebook or Twitter to communicate with other interested travelers. You can add this functionality to an existing widget using ADTs: launch a template editor, create a custom template, and configure your app to host that template. ADTs let you re-skin your application and give you ultimate control over its appearance and functionality.

Creating an ADT

Here's the process of creating an ADT:

1. From Site Administration, click the *Site Selector* button () to choose the Site where you want to create the ADT.

2. Open *Build* → *Application Display Templates*.

If you selected the Global context, this page shows a list of sample templates available for your apps. These sample templates differ from the default templates already configured in the apps. If you choose a Site to host your template, you must create a custom template for that Site's apps.

3. Click the *Add* () button, and you're prompted to select the type of template to create.

- *Asset Publisher*
- *Blogs*
- *Breadcrumb*
- *Categories Navigation*

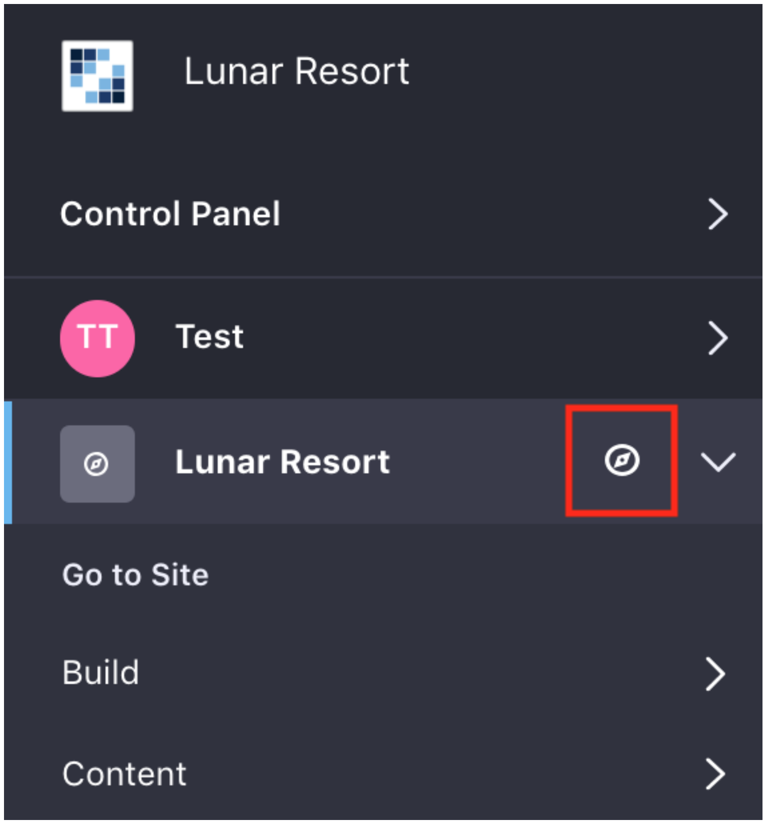

Figure 15.1: The Site Administration dropdown menu lets you choose the context in which your application display template resides.

- *Language Selector*
- *Media Gallery*
- *Navigation Menu*
- *RSS Publisher*
- *Site Map*
- *Tags Navigation*
- *Wiki*

4. Enter the name and, optionally, open *Details* to provide a description and a small image to use. You can select the language type for your template.

5. Within *Details* select a scripting language to use. You can use Freemarker or Velocity. Freemarker is recommended.

6. Use the *Script* section to create the ADT.

7. Click *Save* when done.

The Template Editor

On the left side of the template editor is a palette of common variables used for making templates. This is a great reference when creating your template. To place one of the variables into the template editor, position your text cursor where you want it placed, and click the variable name.

Each variable also has a tooltip which displays a detailed description. Because there are multiple kinds of ADTs, there are also different variables for each ADT. Thus, each template has a different set of variables only applicable for that specific template.

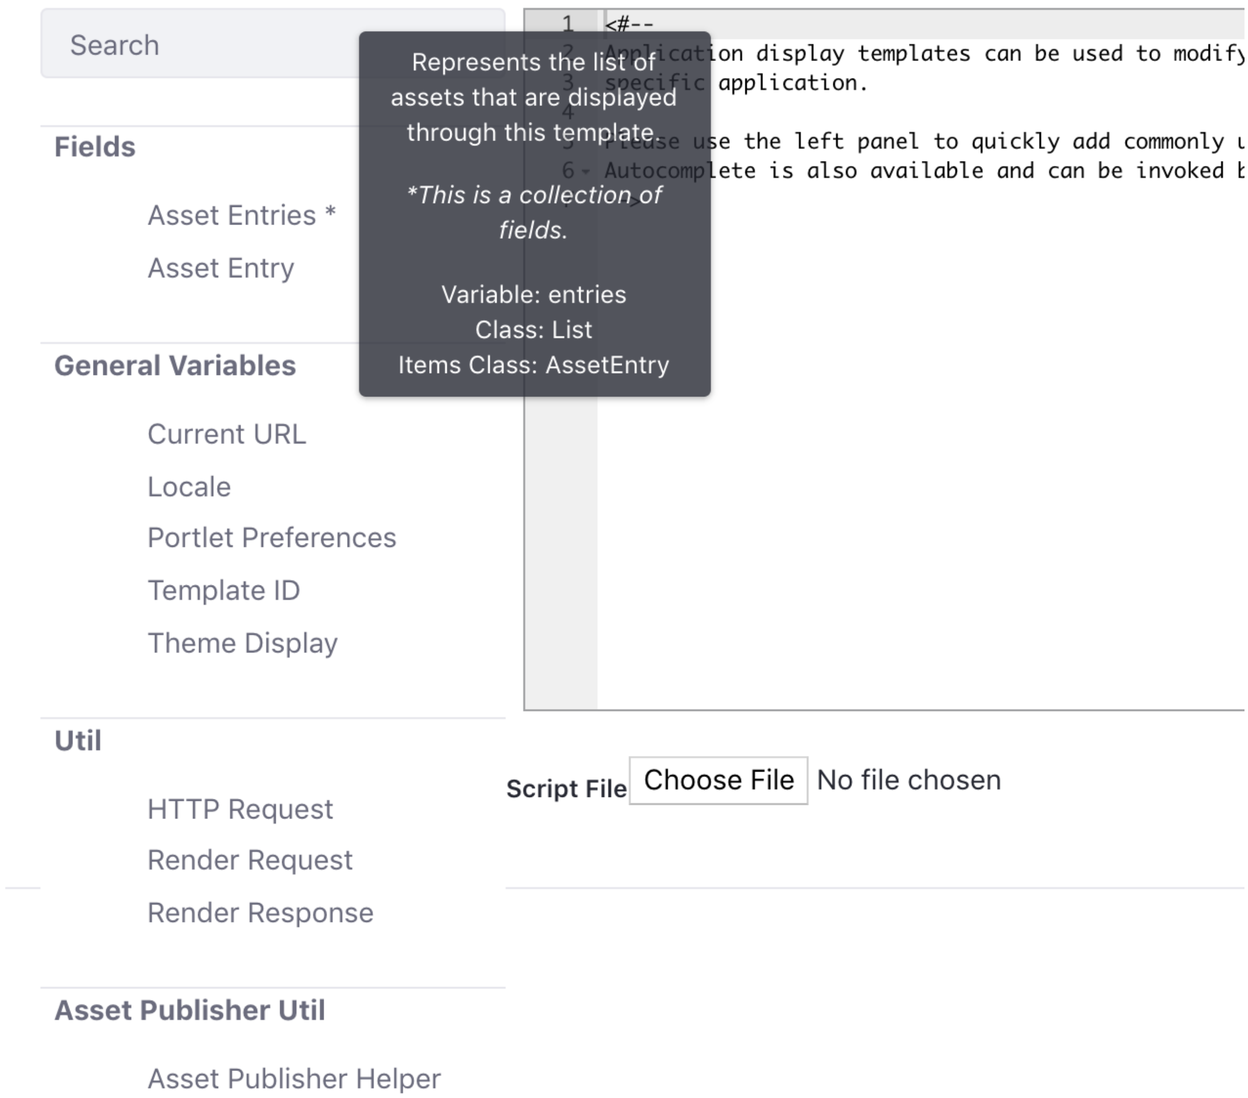

Figure 15.2: Liferay offers a versatile script editor to customize your ADT.

You can also use the autocomplete feature to add variables to your template. It can be invoked by typing *${* which opens a drop-down menu of available variables. By clicking one of the variables, the editor inserts the variable into the editor.

You can also embed same-type templates into other templates. For example, suppose you have an existing Wiki ADT and would like to create another similar Wiki ADT. Instead of starting from scratch, you can import the existing Wiki ADT into your new one and build off of it. In other words, you can utilize ADTs as generic templates which allow for reusable code to be imported by Velocity or FreeMarker templates in the system. For more information on how to create a custom template, visit the Styling Widgets with Application Display Templates tutorial.

Configuring ADTs

After you've saved your ADT, you can manage it through its *Actions* (⋮) button. This provides several options:

- *Edit*: lets you modify the ADT's setup properties.
- *Permissions*: lets you manage the permissions *Update*, *Permissions*, *Delete*, and *View* for the ADT.
- *Copy*: creates a copy of the ADT.
- *Delete*: deletes the ADT.

Additionally, your ADT generates a static URL and a WebDAV URL. These values access the XML source of your template. You can find these URLs by clicking the ADT from the menu and expanding the *Details* section. With the WebDAV URL, Site administrators can add, browse, edit, and delete ADTs on a remote server. If you want to learn more about what the WebDAV URL can do, visit the tutorial on WebDAV access.

Note: Embedding widgets into ADTs, although possible, is not recommended because this could cause conflicts with other widgets or unexpected behavior (e.g., embedding a widget that aggregates data to the breadcrumb). If embedding a widget into an ADT is your only option, make sure it does not interfere with other widgets.

Next you must configure the widget to use the new ADT:

1. Go to the *Configuration* page for the application you want to modify and open its *Display Settings*.

2. Under *Display Template*, select your ADT from the drop-down menu.

Also, you can manage Site-specific display templates for your app: do this by clicking the *Manage Display Templates for [SPECIFIC_SITE]* link next to the *Display Template* drop-down menu. A window appears with a list of your configured templates only available for your Site with options to add new templates or edit existing templates.

15.2 ADT Example

Now that you know the general functions of ADTs, you'll create your own. This brief demonstration will show you just how easy, yet powerful, ADTs can be for your Liferay instance.

1. Add the Media Gallery widget to a page by navigating to *Add* (+) → *Widgets* → *Content Management* → *Media Gallery*.

2. Click the widgets's *Add* button (+) → *Multiple Media* and select two custom photos to display. Then click *Save*, and navigate back to the main application screen.

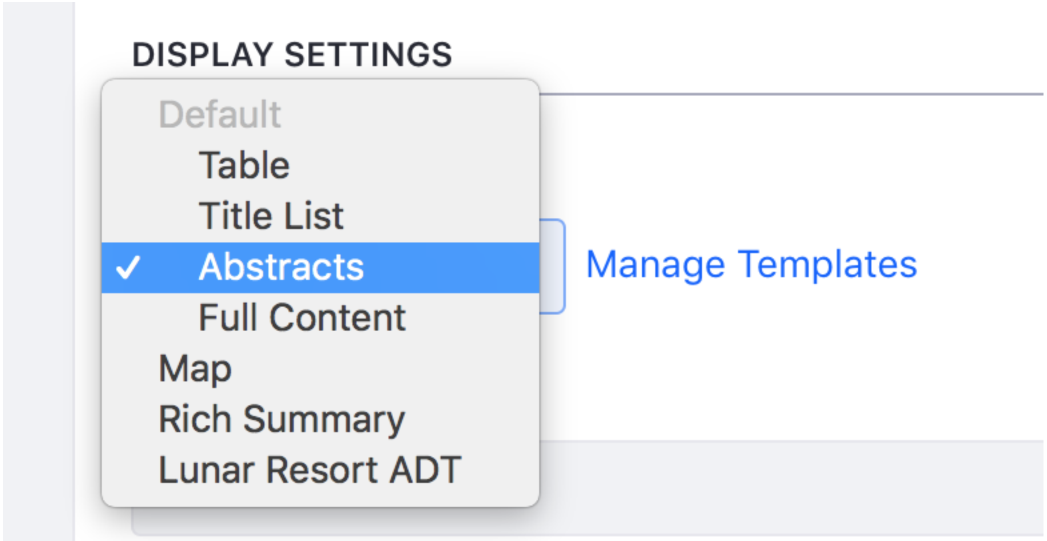

Figure 15.3: In the *Configuration* menu of an app, you can edit and manage available ADTs.

3. Notice the default format of the pictures. To change the display template for this widget, navigate to *Options* (⋮) → *Configuration*.

4. From the *Display Template* drop-down menu, select *Carousel*. Then click *Save*.

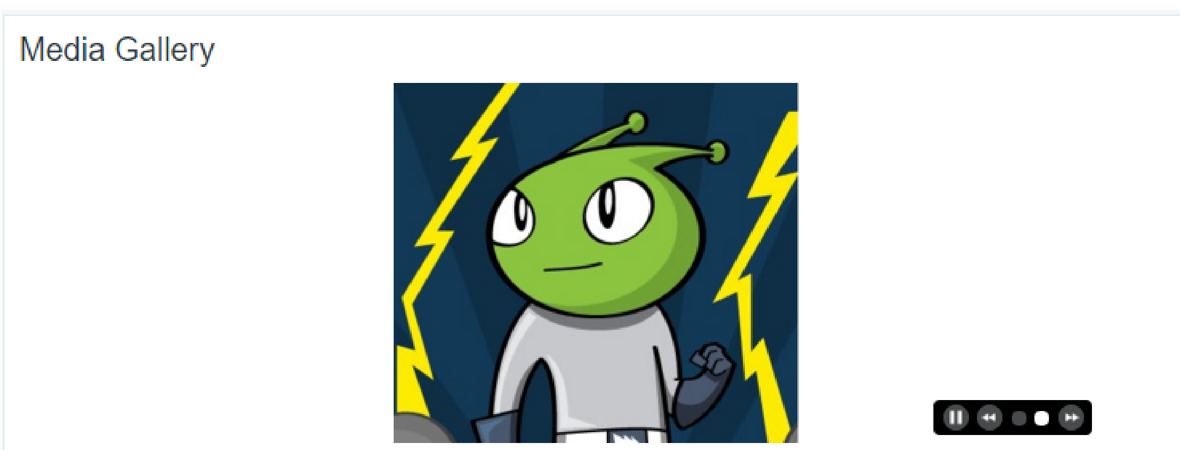

Figure 15.4: After applying the Carousel ADT, your pictures are displayed as a carousel slideshow.

The Media Gallery application is transformed into a carousel slideshow. At this time, it's perfectly natural to be experiencing "I can conquer the world" feelings, just as Liferay's mascot, Ray, exudes in the image above. ADTs have that kind of power to transform your site into an enjoyable and convenient home for users.

Customizing the user interface of Liferay DXP's bundled widgets provides the ultimate customization experience for Liferay users.

15.3 Setting a Default ADT

You can change the ADT for an individual widget through its own configuration, but to configure the default ADT for all widgets of that type, you must go to **System Settings**. In the System Settings you can find a configuration for every widget in Liferay. Any widget that supports ADTs has a *Display Style Group ID* and a *Display Style* option.

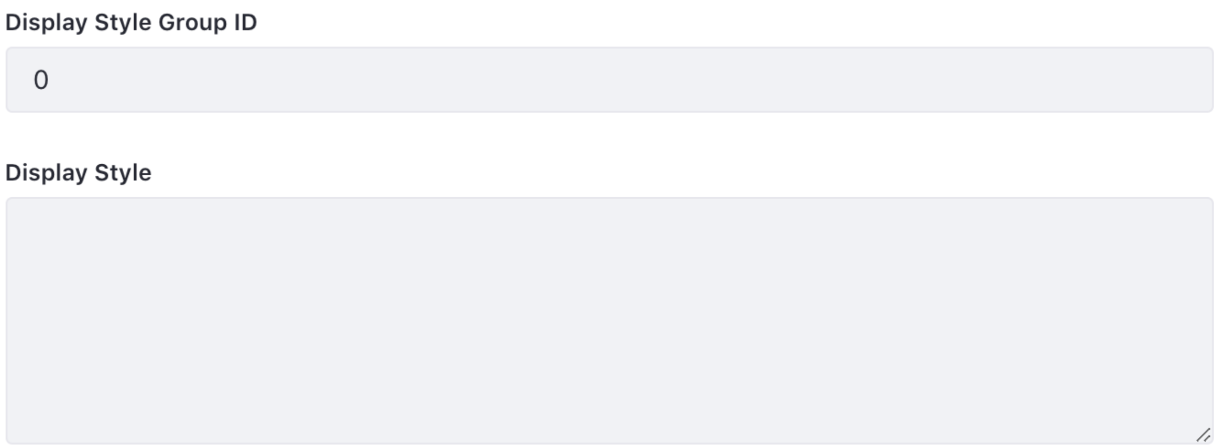

Figure 15.5: The ADT configuration in System Settings lets you change the display style.

- **Display Style Group ID:** The Site ID where the ADT is located. For Global templates use *0* for the ID.
- **Display Style:** The ADT's template key.

To enter a Display Style, you first need the **Template Key** for the template you want to use. To get the Template Key, go to the *Application Display Template* list for a given Site and retrieve it from the ADT listing. Then enter the display style as `ddmTemplate_[template-key]`.

Default ADT Example

For example, configure the Language Selector ADTs like this:

1. Open the *Menu*.
2. Using the *Site Selector*, select the *Global* site.
3. Go to *Build* → *Application Display Templates*
4. Find the *Icon Language Selector Template* and click on it.
5. Open *Details* and find the *Template Key* - `LANGUAGE-ICON-FTL`

Now that you have the ID, you can change the template from System Settings.

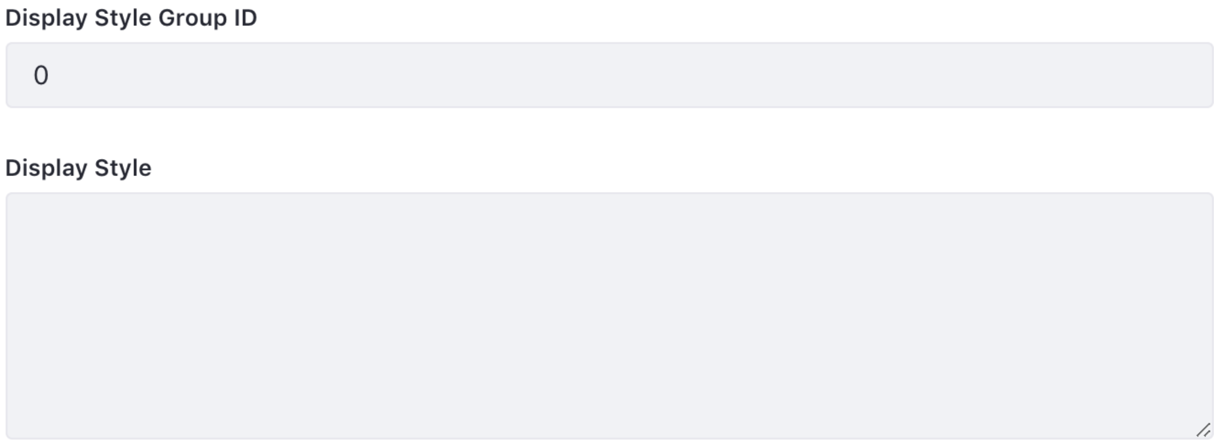

Figure 15.6: System Settings shows where you can find the Template Key.

1. Once you have the ID, go to the *Control Panel* → *System Settings*
2. Find *Localization* under the *Platform* heading and select *Language Selector* from the options on the left.
3. In the *Display Style* field, enter `ddmTemplate_LANGUAGE-ICON-FTL`.

Now any Language Selector widgets are added to a page use the new defaults. This doesn't affect widgets already added to a page and configured.

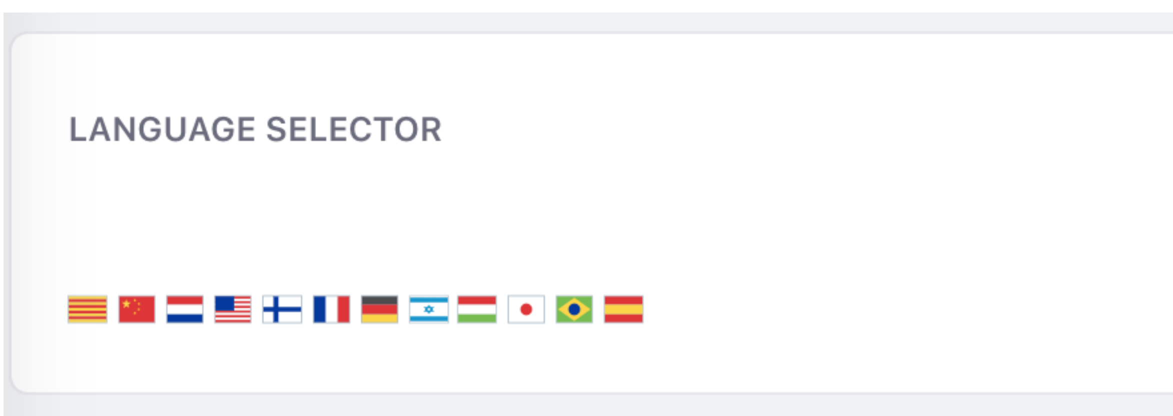

Figure 15.7: You can see the new default configuration.

Chapter 16

Customizing Page Options

Every page has options to give it a unique configuration. You can handle that configuration at the individual page level or configure a Page Set of public or private pages all at once.

When you configure options across Sites, Page Sets, and individual pages, there is a hierarchy that flows down. For example, you can set a theme at the Site level that applies to all pages within a Site. You could then set a different theme for the Private Pages set. You can even set a different theme for a specific page that is different than the master site configuration or the configuration for its page set.

When configuring pages, be aware of your context, so you don't configure an option that you intended for the whole site just on one page—or change the whole site's configuration, when you meant to change it for only a specific page.

Chapter 17

Configuring Page Sets

To configure options for the entire Page Set, select the *Configure* icon next to the Page Set in *Pages*. Options configured for the Page Set apply to all its pages. Page Set options override options set at the Site level, and customizations to an individual page override those for the Page Set.

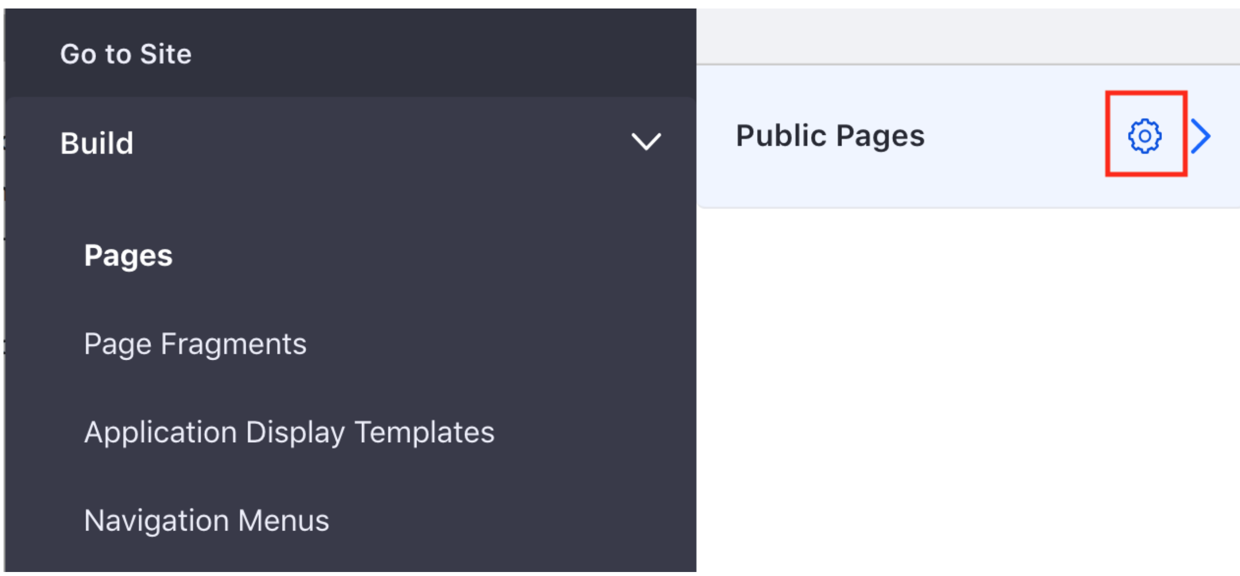

Figure 17.1: Selecting the Page Set configuration option.

17.1 Page Set Look and Feel

Page Set configuration starts with the *Look and Feel* tab. Here you have an interface for choosing a theme for the current Site.

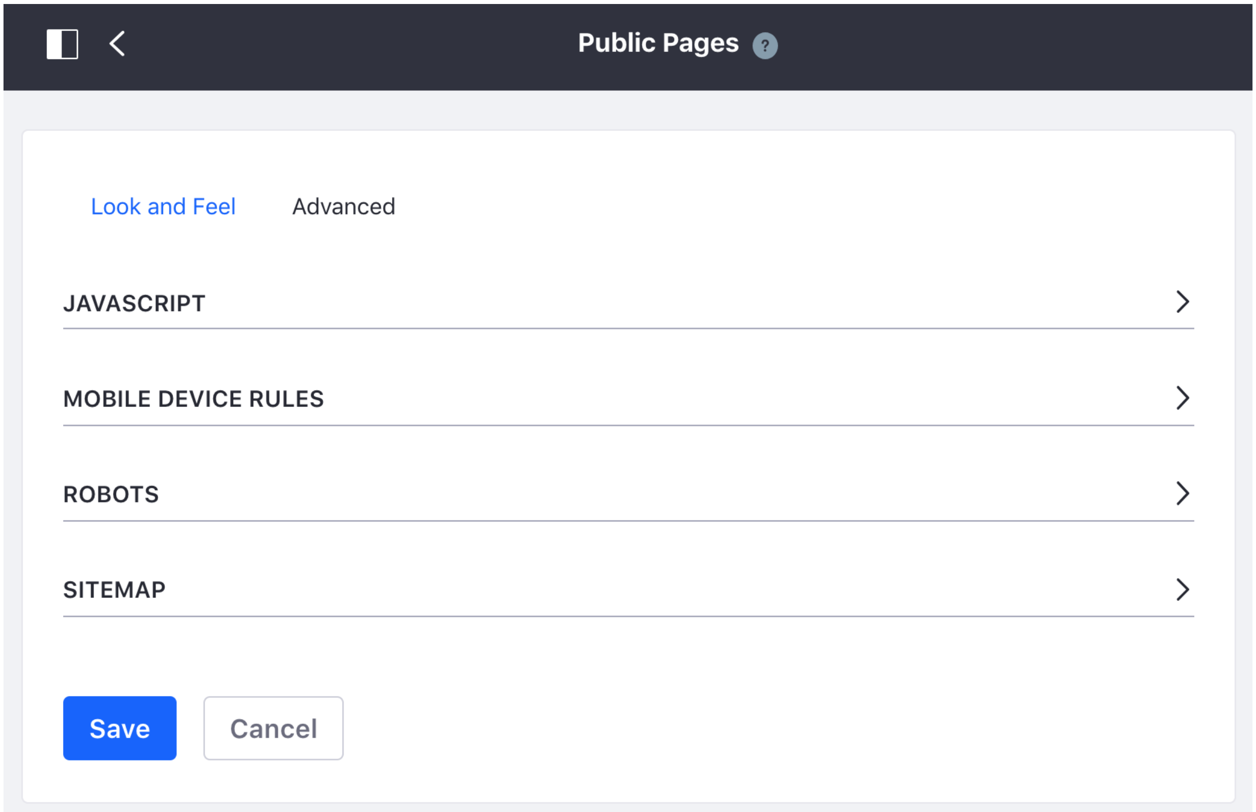

Figure 17.2: The Look and Feel page set tab.

Themes

Themes can transform the entire look of your Site. They are created by developers and are easily installed using Liferay Marketplace.

You can apply themes to the entire Site or to individual pages. For the Site, go to *Pages* → the Site (public or private), and click the Gear icon. For individual pages, click *Configure* → *Define a specific look and feel for this page* option under the page's *Look and Feel* category. See the Themes and Layout Templates tutorials section for information on creating and developing your own custom themes.

Many themes include more than one color scheme, which keeps the existing look and feel while giving the Site a different flavor. The Color Schemes option is not available for the default theme.

There are a few more configurable settings for your Page Set look and feel. You can switch the bullet style between dots and arrows and you can choose whether or not to show maximize/minimize application links by default. The *CSS* section lets you enter custom CSS for tweaking your theme.

Using a Custom Logo for a Site

By default, the Liferay logo is used for your Site pages' logo. If you want to use your own logo for a specific Site, select the *Logo* tab from the *Configure* interface and browse to the location of your logo. Make sure your logo fits the space in the top left corner of the theme you're using for your website. If you don't, you could wind up with a Site that's difficult to navigate, as other page elements are pushed aside to make way for the logo.

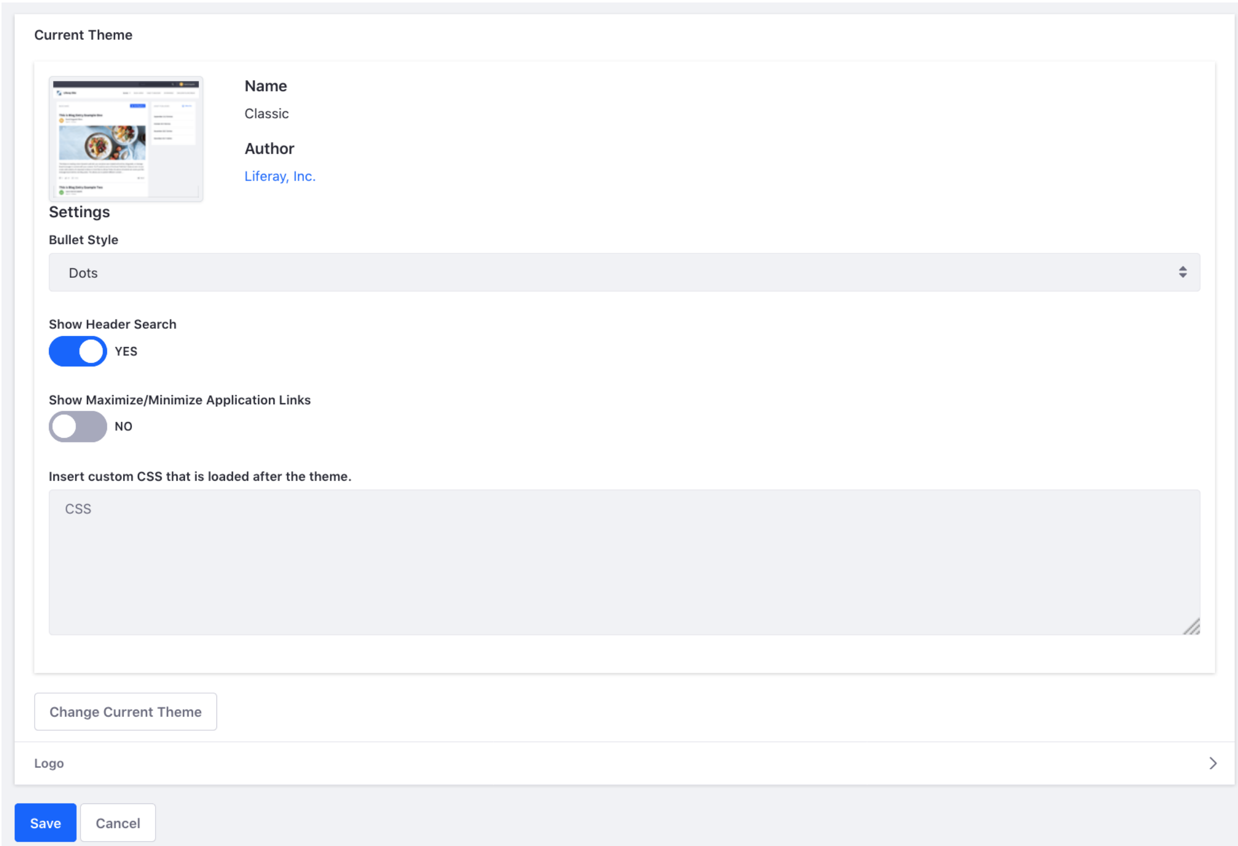

Figure 17.3: The Look and Feel interface allows you to choose a theme for the current site.

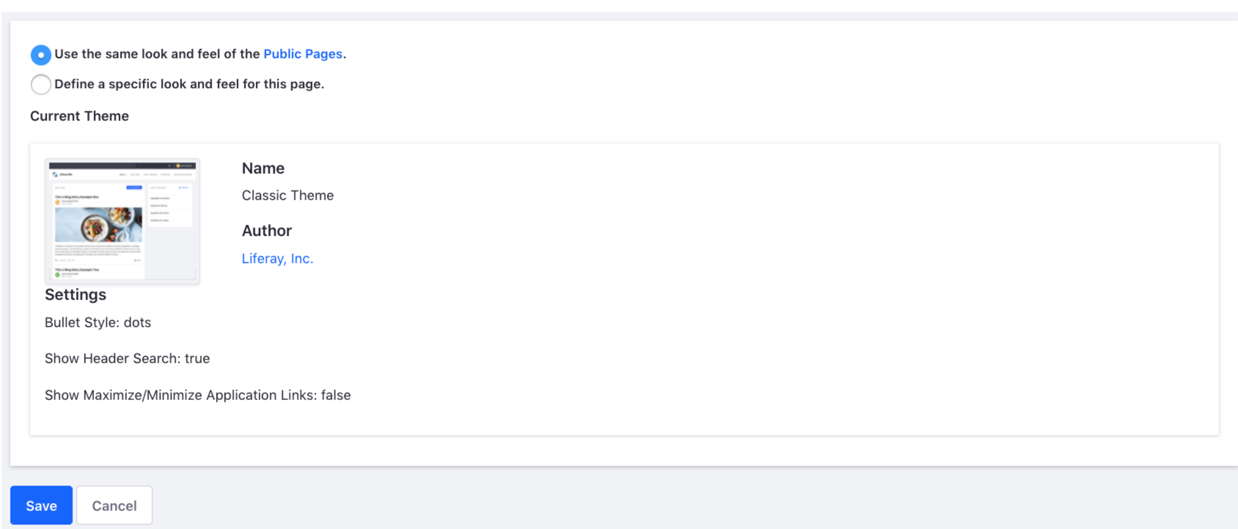

Figure 17.4: You can define a specific look and feel for a page.

In the logo tab, you can also choose whether or not to display the Site name on the Site. If you check the box labeled *Show Site Name*, the Site name appears next to the logo. This option is enabled by default and cannot be disabled if the *Allow Site Administrators to set their own logo* option is disabled in *Instance Settings*. Removing the Site name is not available for the default Site—you can configure this only for new Sites and user pages.

17.2 Advanced Page Set Options

There are some powerful options that should only be used by those with a firm command of the technology, or they could have major unintended side effects. You can find these options under the *Advanced* tab.

Executing JavaScript in Site Pages

At the top of the *Advanced* tab is a JavaScript editor. Code entered here is executed at the bottom of every page in the Site. If your Site's theme uses JavaScript (as is usually the case), it's best to add custom JavaScript code to the theme and *not* here. This way, all your Site's JavaScript code remains in one place.

This may be useful if your Site's theme does *not* use JavaScript. In this case, you can place *all* of your Site's JavaScript here.

Rendering Pages for Mobile Devices

Mobile Device Rules lets you configure your page set to have specific behaviors for specific mobile devices or types. Mobile device rules are inherited from your Public Pages, but you can define specific rules per page. You can edit the Look and Feel of specific pages for mobile devices, including the theme. This is explained in Mobile Device Rules.

Robots

The *Robots* option lets you configure `robots.txt` rules for the domain: both its public and private pages. The `robots.txt` file provides instructions to search engines and other tools that are automatically crawling and indexing your Site. Common entries here include defining some pages not to be indexed.

Notifying Search Engines of Site Pages

The *Sitemap* option generates a sitemap you can send to some search engines so they can crawl your Site. It uses the industry standard sitemap protocol.

Select a search engine link to send the sitemap to it. It's only necessary to do this once per Site.

If you're interested in seeing what is sent to the search engines, select the *preview* link to see the generated XML.

Chapter 18

Configuring Individual Pages

After you've configured your Page Set, you can reconfigure some options at the individual page level. When you decide to customize a single page, options that were not available when initially creating a page appear. You can customize a page by navigating to *Pages* under *Build* menu and selecting *Options* (⋮) → *Configure* next to the page you want to edit from the navigation tree. Alternatively, you can click the *Configure* icon on the top right of any page.

18.1 Individual Page Settings

On the Configure page are four tabs: General, SEO, Look and Feel, and Advanced. Options selected here have no effect on the rest of the Site; just the page you've selected. Many of these options are the same as those that configure the complete page set, so you can view more details in the Configuring Page Sets article.

Note that many of the options are localizable, so you can provide translations based on the user's locale.

General

The *General* tab lets you configure the basic information and design for the page. You can change the *Name*, *Friendly URL*, and *Page Layout*.

Name and Friendly URL

The *Name* is the title that appears in the browser's title bar, and how the page is identified in the navigation. The *Friendly URL* defines the page's link. It is a best practice to have the URL match the name of the Page, so these two should generally be updated together.

Page Layout

For Widget Pages, you can select a Layout Template that defines droppable locations for widgets. Layout templates define a number of sections with columns and rows. Widgets added to a section expand (or contract) horizontally to fill the space and can be stacked vertically.

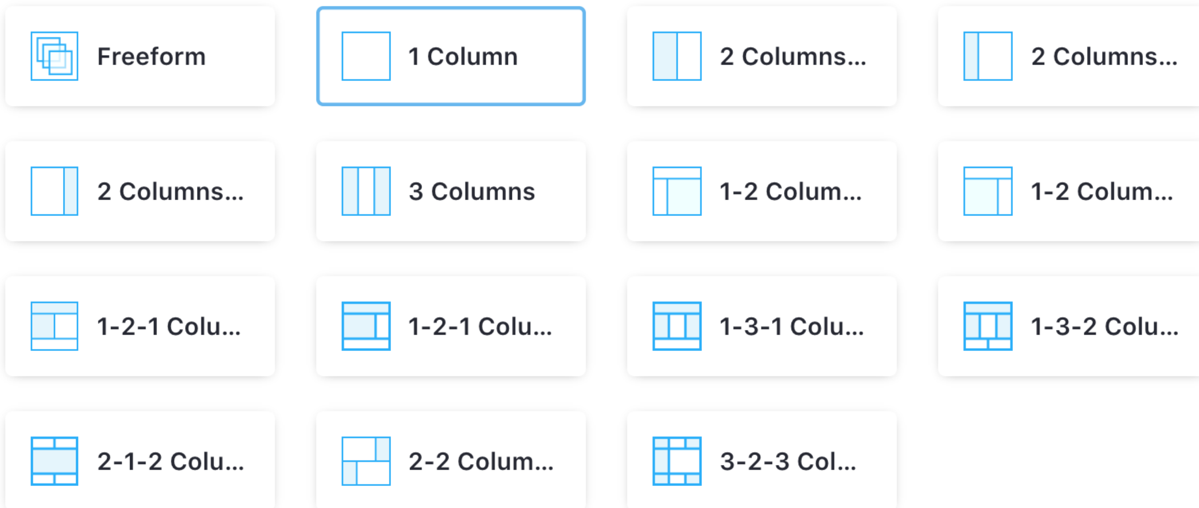

Figure 18.1: Setting a layout template for your page.

Categorization and SEO

Managing your page's content drastically improves your page's organization and user experience. The Site page's configuration options offers some opportunities to organize page content.

Categories

The *Categories* tab shows the categorization options. These tools help administrators organize the page so users can find your page and its content through search and navigation. For more information on using tags and categories, see Organizing Content with Tags and Categories.

SEO

SEO provides several ways to optimize the data the page provides to an indexer that's crawling the page. You can set the various meta tags for description, keywords and robots. There's also a separate Robots section for telling robots how frequently the page is updated and how it should be prioritized. If the page is localized, you can select a box to generate canonical links by language. If you want to set some of these settings for the entire Site, you can specify them from the Sitemaps and Robots tabs of the Manage Site Settings dialog box (see below).

Each asset (web content article, blog entry, etc.) has a unique URL. From the search engine's point of view, this makes your pages rank higher since any references to variations of a specific URL are considered references to the same page.

Look and Feel

Look and Feel lets you set a page-specific theme. You can inherit what you already have configured for your Page Set's theme, or you can define a theme per page. See Customizing the Look and Feel of Site Pages for more details.

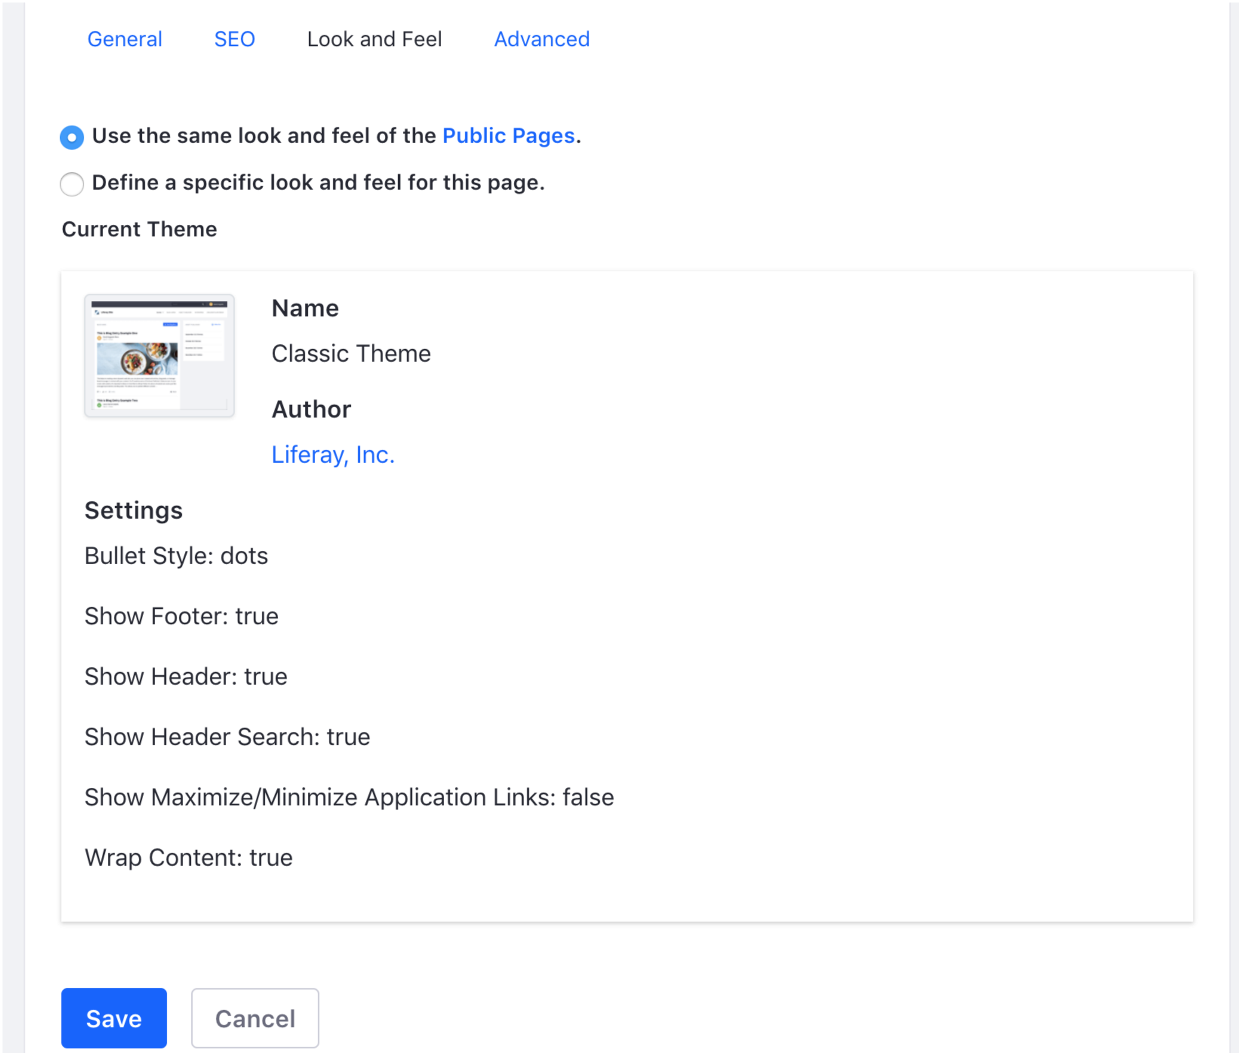

Figure 18.2: Viewing the Look and Feel page configuration. .

Advanced Settings

Advanced Settings contains options useful for specific cases. Some of these are the same as the options available at the Site or Page Set level, but *Custom Fields, Embedded Widgets,* and *Customization Settings* are unique to the individual page configuration.

Query String

You can set a query string to provide parameters to the page. This can become useful to web content templates. You can set a target for the page so that it pops up in a particularly named window or appears in a frameset. And you can set an icon for the page that appears in the navigation menu.

Custom Fields

Custom Fields lets you edit the custom fields you already have configured for the *Page* resource. If you don't have any custom fields configured in your Site, this option doesn't appear. In this case, navigate to the Control Panel → *Custom Fields* located under the *Configuration* tab. These are metadata about the page and can be anything you like, such as author or creation date. For more information on Custom Fields, see Custom Fields.

Mobile Device Rules

Apply rules for how this page should render for various mobile devices here. Create them by navigating to Site Administration menu and selecting *Configuration* → *Mobile Device Families*.

Embedded Widgets

This option only appears if you have embedded one or more widgets on the page.

Widgets can be embedded on a page via web content template or fragment. To learn more about this, see Adding Templates. You can embed a widget on a page layout or theme programmatically. If you're interested in learning more about this, visit the Embedding Portlets in Themes tutorial.

Customization Settings

This configuration option in the *Advanced* tab lets you mark specific sections of the page you want users to be able to customize. You can learn more about page customizations in Personalizing Pages.

JavaScript

This shows a JavaScript editor for code that's executed at the bottom of your page. If your Site's theme uses JavaScript (as is usually the case), it's best to add custom JavaScript code to the theme instead. This way, all your Site's JavaScript code remains in one place.

This configuration option is also available for Page Sets like Public Pages and Private Pages. Visit Executing JavaScript in Site Pages for more information on doing this for Page Sets.

18.2 Personalizing Pages

Administrators can designate pages or sections of Widget Pages as customizable. When a user visits such a page, a notification appears stating that the user can customize the page. Users can make customizations only in the sections of pages designated by administrators. Customizations are based on the rows and columns of a page layout. Page customizations are only visible to the user who made the customizations. By default, Site members can make page customizations but non-Site members and guests can't.

Enabling Page Customizations

To enable page customizations as an administrator,

1. Click *Configure Page* from the *Options* button next to the Page you want to let Site members modify.

2. Select the *Advanced* tab at the top of the page and expand the *Customization Settings* area.
3. Click the *Customizable* selector button to activate customizations.

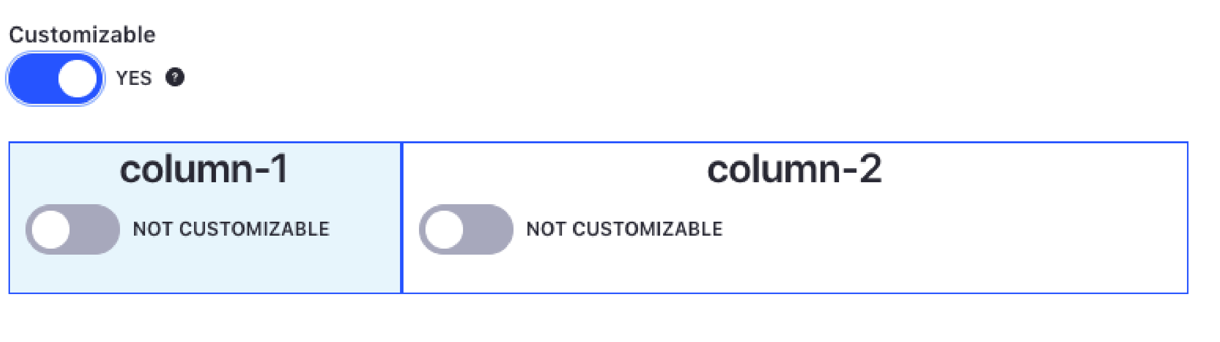

Figure 18.3: To enable page customizations, click on the *Configure Page* button next to the page, expand the *Customization Settings* area, and click on the *Customizable* button.

4. Select the sections of the page that should be customizable.
5. Enable one or more of the *Customizable* sections so Site members can customize sections of the page. Regions that you've designated as customizable are colored blue.

When Site members visit your customizable page, they see an extended Control Menu with a notification saying *You can customize this page*. Site members can toggle whether to view or hide the customizable regions. If you toggle the selector to view customizable regions, the regions on the page are color-coded to help distinguish customizable vs. non-customizable sections of the page.

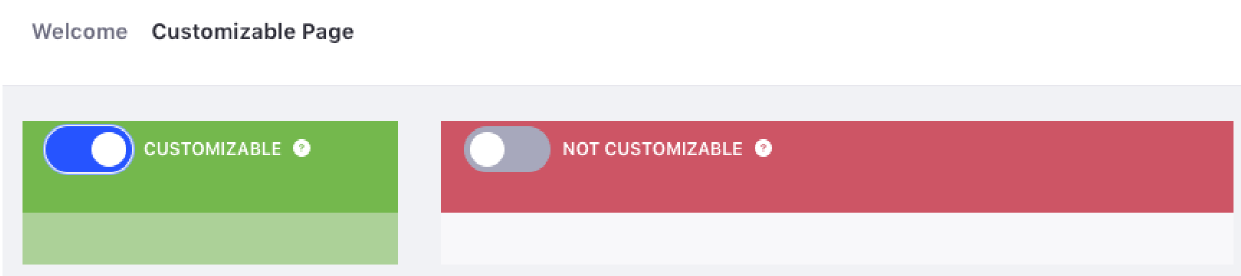

Figure 18.4: Customizable regions are colored green and non-customizable regions are colored red.

Customization Permissions

Administrators must grant users permission to customize pages under the Site section. This can be achieved by assigning permission to a Role, then assigning this Role to the appropriate users. For example, if you want users to be able to customize your customizable pages, assign the *Customize* permission to the Role *User*. If you want Site members to be able to customize their Sites' customizable pages, accept the default setting. By default, the *Customize* permission is assigned to the Role *Site Member*.

The *Customize* permission also lets users customize the look and feel of apps and import or export app settings.

Customizing Pages

With customization active, Site members can access the Add menu from the top right side of the screen when viewing their customizable page, which lets them add apps to the customizable sections of the page. If they click *View Page without my customizations*, the Add menu disappears.

Users can make two kinds of customizations to customizable regions:

1. They can configure applications within the customizable regions.

2. They can add apps to or remove apps from the customizable regions.

Reset My Customizations from the *Options* button restores a user's customized page to match the default page, discarding their customizations so they can start anew.

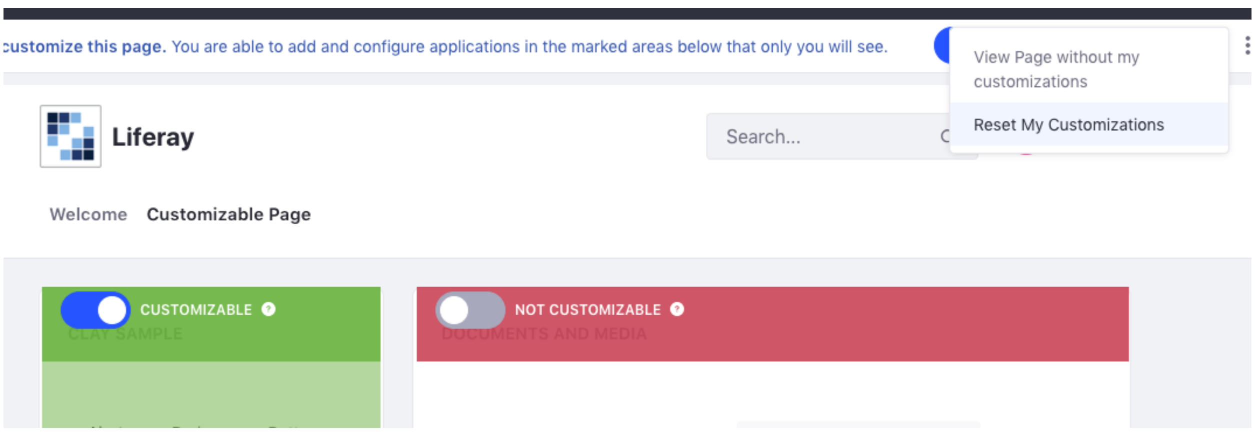

Figure 18.5: Customizable areas are highlighted green when organizing apps on the page.

Users can't change a non-instanceable app's configuration inside a customizable region since those apps are tied to the Site where they've been added.

Viewing Customized Pages

Site members can also choose between viewing their customized page and viewing the default page by selecting the *Options* button () from the Control Menu and clicking the *View Page without my customizations* or *View My Customized Page*.

Administrators of customizable pages have the same two views as Site members: the *default page* view and the *customized page* view. Changes made to the default page affect all users, whereas changes made to the customized page affect only the administrator who made the changes. Changes made by administrators to non-customizable sections in the default view are immediately applied for all users. Changes made by administrators to customizable sections, however, do *not* overwrite users' customizations.

Customization Example

As an administrator,

1. Go to *Page Configuration* for the *Welcome* page.
2. Go to the *Advanced* tab, and activate Customizations.
3. Set the main column of the Welcome page of the Lunar Resort Site to be customizable.

As a regular user,

1. Navigate to the *Welcome* page.
2. Click *Add* → *Widgets*.
3. Locate the *Language Selector* widget and add it to the page.

The Language Selector application lets users select their language to view a translation of your Site into their native language. After closing the Configuration dialog box of the Language Selector app, the customized Welcome page looks like this:

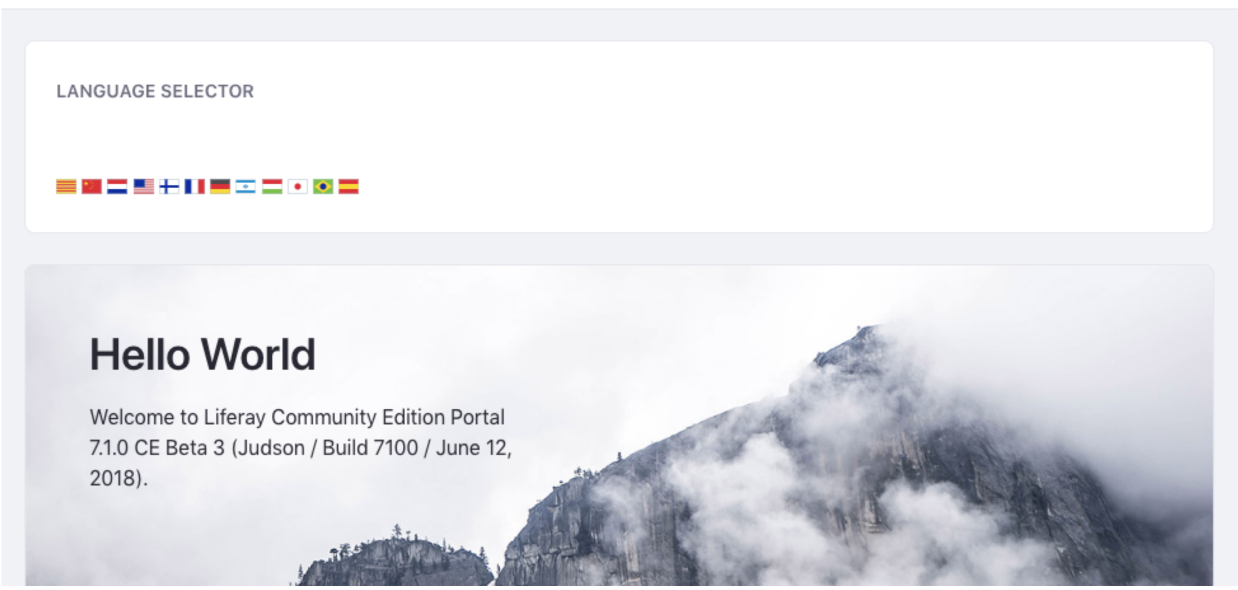

Figure 18.6: In this example, the user added the Language app, and changed the display style from icons to a select box.

18.3 Changing Page Permissions

Public pages are just that: public. They can be viewed by anybody, logged in or not. And private pages are only private from non-members of the Site. If someone has joined your Site or is a member of your organization, that person can see all the private pages. If you want to further protect some content, you can modify the permissions on individual pages in either page group so only certain users can view them.

Here's how to create a page only administrators can see:

1. Go to your Site's Site Administration dropdown and select *Build* → *Pages* → *Private Pages*.

2. Create a page called *Admin Tips*.

3. Click *Configure Page* from the Options button dropdown for the page in the left menu.

4. Select *Permissions* from the *Options* icon () in the top right corner of the screen.

5. Uncheck the *View* and *Add Discussion* permissions next to the Site Member role.

6. Click the *Save* button.

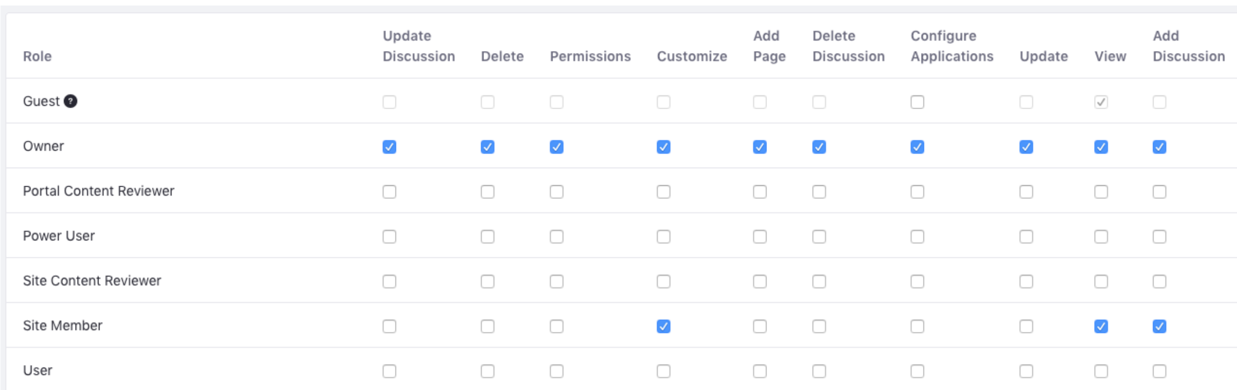

Figure 18.7: The Permissions offer a plethora of options for each role.

Congratulations! You've changed the permissions for this page so only Site administrators can view it. Any users you add to this Role can now see the page. Other users, even members of this Site, don't have permission to see it.

Pages are as flexible as pages you'd create manually without Liferay. Using a point and click interface, you can define your Site any way you want. You can create and remove pages, export and import them, set their layouts, define how they are indexed by search engines, and more.

CHAPTER 19

MANAGING MEMBERS IN YOUR SITE

Users and Sites are important concepts. Sites are where all your content and pages are stored, and Users access and create that content. While user management is covered in depth in our User Management tutorial, there are other user configuration options specific to Site Management:

- Adding members to Sites administratively
- Adding members to Sites automatically
- Creating Teams of Site members for various functions

This section of tutorials shows you how to manage and configure these Site options for Users.

19.1 Adding Members to Sites

In Adding Sites you learned the difference between Site Membership Types and about public and private pages within a Site. Now you'll learn how to add users manually to Sites and how to provide options for self management. For review, there are a few key reasons why Site Membership management is important:

1. Only Site Members can view a Site's Private Pages.
2. Site Members have more permissions than guests for many widgets like Message Boards and Wikis that enable them to create content and collaborate on your Site.
3. Site Members can be associated with Roles that grant Site privileges.

Administrating Site Membership

Administrators can manage Site members from that Site's *Site Membership* page.

1. Open *Site Administration* and select the Site that you want to manage members for.
2. Click on *Members → Site Memberships*

From here you can manage Site Memberships, Organization, and User Group associations. You can learn more about those in the Users and Organizations tutorial. Here you see a list of all of the current Users of the Site and you can add or remove user memberships from the Site.

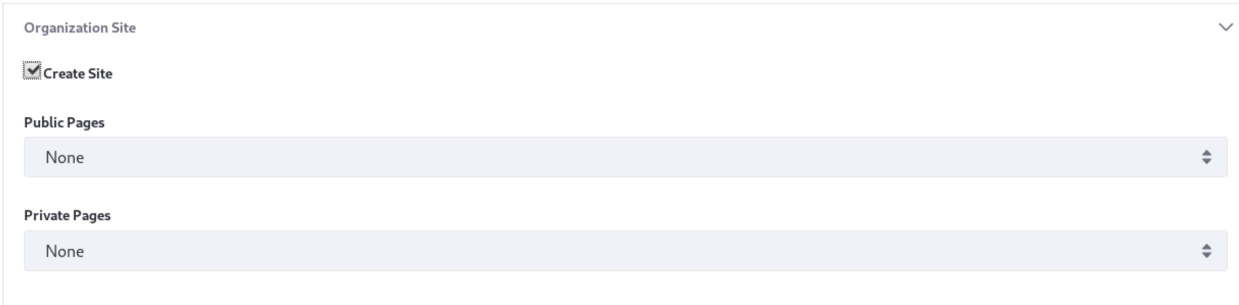

Figure 19.1: The current members of the Site as displayed on the *Site Memberships* page.

Adding Members to a Site

Follow these steps to make an existing user a member of the Site:

1. Click the *New* (+) button in the top right of the screen.
2. Use *Filter and Order* or the *Search* function to locate the User you want to add to the Site.
3. Select the User(s) you wish and click *Done*.

On the *Assign Users to This Site* screen, all Users eligible to be added to the Site appear. Deactivated Users do not appear. Site members also appear, but with a greyed-out checkbox.

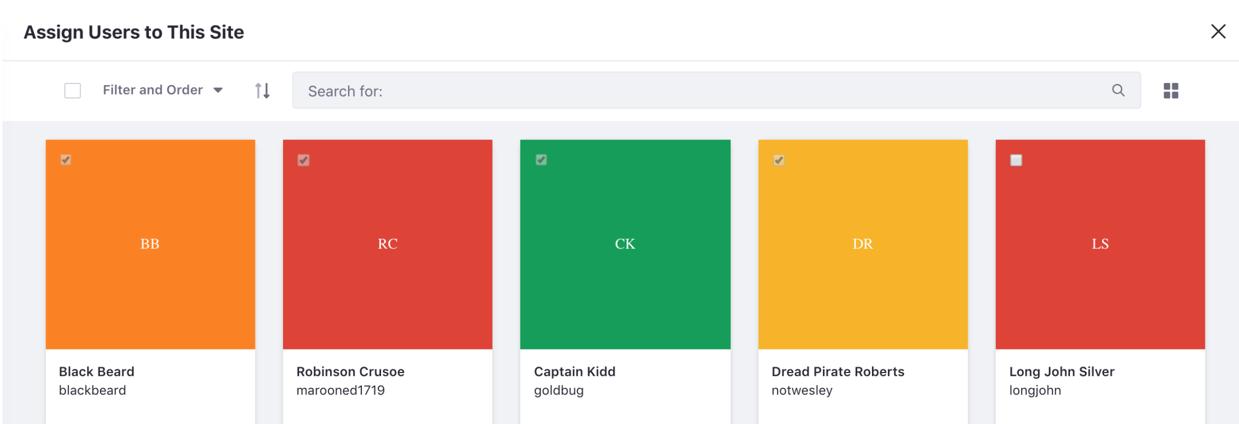

Figure 19.2: The list of users available to add to the current Site. Note that the current members are visible but cannot be added or removed here.

Removing User Membership from a Site

There are two ways to remove a user from a Site. You can remove an individual member like this:

1. Click the *Actions* (⋮) icon for the user that you want to remove.
2. Select *Remove Membership*.

3. In the pop-up that appears, confirm the removal.

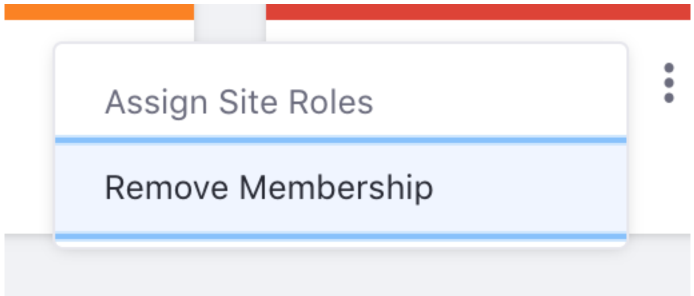

Figure 19.3: Selecting to remove a user.

To remove several users at once, you can do this:

1. Click the checkbox for each user that you want to remove.
2. In the menu at the top of the page, click the trash can icon to remove the users from the Site.
3. In the pop-up that appears, confirm the removal.

Removed Users lose access to the Site's private pages and membership in any Site Roles or Teams they had.

Assigning Site Roles

Roles grant permissions in Liferay DXP. Roles can be assigned for the entire instance or just for one specific Site or Organization. Site Roles assign permissions for a specific Site.

You can use the same interface options that you used to remove Users from the Site to assign them to Site Roles. If you select a User or Users and click *Assign Site Roles* (either through the Actions menu or the menu at the top), you are taken to the *Assign Site Roles* screen. From here:

1. Select the Roles that you want to assign to the selected users.
2. Click *Done*.

Site Roles are created at a global level, but when they're assigned they only provide privileges for the specific Site where they were assigned. Since Roles are created at a global level, they cannot be created by Site Administrators (since Site Administrators only have Administrator privileges for their Site). **Teams** allow Site Administrators to assign permissions to groups of Users within their Sites. Next, you'll look at more configuration for managing members of your Site.

19.2 Managing Site Settings for Members

Now you'll learn how to set some global settings for your Sites: site membership types, default user associations, and Site ratings.

1. Open the *Product Menu* and use the site selector to choose the site you want to configure.
2. In Site Administration, go to *Configuration* → *Site Settings*.

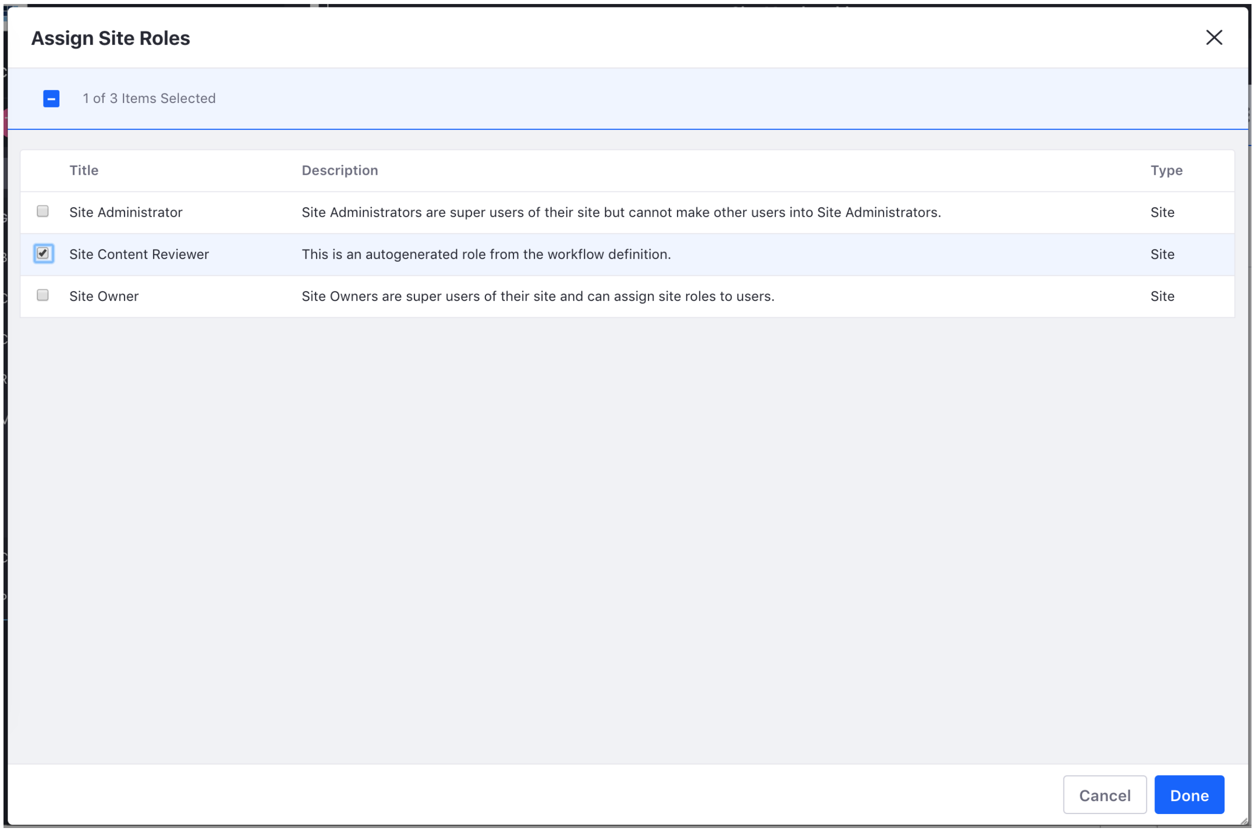

Figure 19.4: Assigning Site Roles.

Details

Details shows the same form you filled out when first creating your Lunar Resort Site. This allows an administrator to change the description and membership type of a Site. The membership type can be set as open, restricted, or private. Users can join and leave an open Site at will. To join a restricted Site, a User must be added by the Site administrator. A User can also request membership through the Sites section of the Control Panel. A private Site is like a restricted Site but doesn't appear in the Sites section of the Control Panel for Users who aren't members.

You can organize Sites into hierarchies. At the bottom of the Details sub-section is the Parent Site section, where you can select the parent site for the currently selected Site. After selecting a parent Site, you have a checkbox option to limit membership to members of the parent Site.

Default User Associations

When you have saved your Site settings, you can begin assigning Users to Roles and Teams. Still inside *Site Settings*, go to the *Advanced* section. *Default User Associations*, the leading option when opening the Advanced tab, lets you configure Site Roles and Teams that newly assigned Site members get by default. To learn more about creating Roles, visit the Roles and Permissions tutorial.

Ratings and Mentions

Next, open the *Social* section of *Site Administration*. From *Ratings* you can enable ratings for applications in the Site. You can also select ratings types like Stars, Likes, and Thumbs.

The last configuration option in Site Settings related to managing Users is *Mentions* in the Social tab, where you can enable/disable Mentioning Users. This lets Users *mention* (notify and/or draw attention to) friends and colleagues by entering the "@" character followed by their user names.

Next, you can learn about Teams.

19.3 Creating Teams to Empower Site Members

If you have an *ad hoc* group of Users who perform the same set of tasks in a Site, you can organize them into Site Teams, and then assign the team permissions for various site-specific functions. Site Teams are the preferred method for collecting permissions within a single Site. Some common functions to assign a Site Team include

- Moderating site Wiki content
- Managing Message Boards threads
- Writing blogs
- Editing a specific page in the site

If your Site has Message Boards, you might want to enable a subset of the Site's members to moderate the categories and threads, and perhaps to ban abusive/offensive posters. To do this, you could create a Site Team named *Lunar Resort Message Board Moderators*, define the team's permissions in the Message Boards application, and assign the desired Site members to the team.

The permissions assigned to a Site Team only apply to that Site. The two key features of Teams are that they are limited to their Sites and that they empower Site Administrators to manage permissions for their Sites since Site Administrators cannot create new Roles.

Note: To create and apply permissions for a group of users to use across multiple Sites or Organizations in your Liferay instance, consider aggregating the Users into a User Group and assigning the User Group permissions via Roles.

To create a team within a Site,

1. Go to the Site Administration page of your Site.

2. Select *Members* → *Site Teams*.

   It's important to note that configuring other Site membership groupings, such as *Users*, *Organizations*, and *User Groups* can be done in the *Site Memberships* app, which is also in the Members tab. You can visit User Management for more information on how Site memberships work.

3. Finally, click the *Add Team* icon ().

4. Enter a name and a description and click *Save*. Your new team shows in the list.

5. To add members, click on the team name link and then select *Add Team Members*.

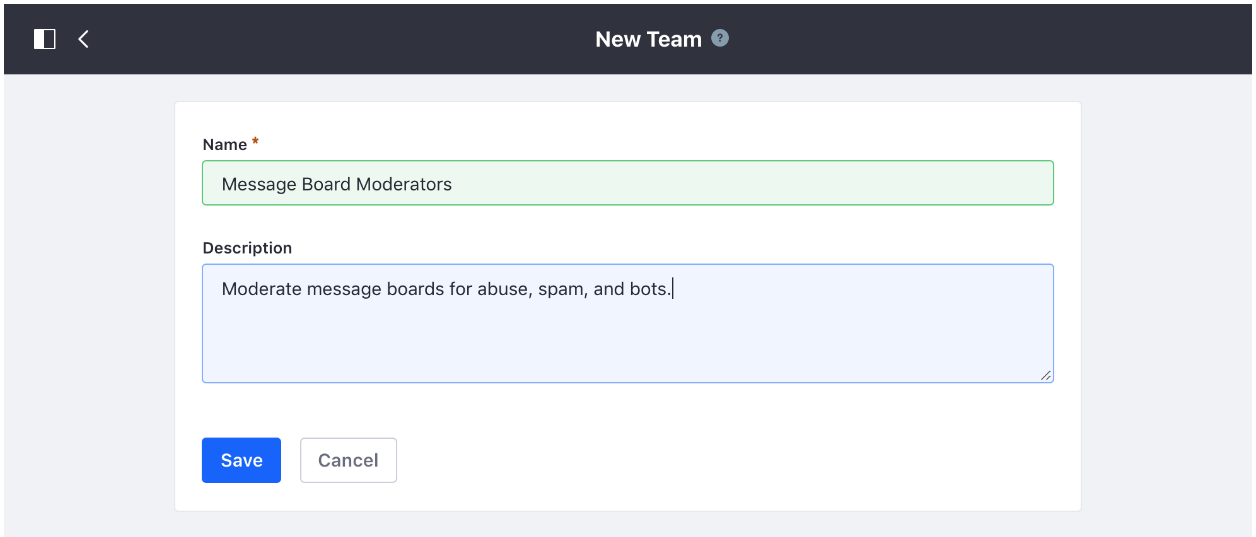

Figure 19.5: Creating teams within your site can foster teamwork and collaboration, as team permissions enable team members to access the same resources and perform the same types of tasks.

To manage a team's permissions, click the *Actions* icon ( ) and select *Permissions* for that team. Setting permissions for the team assigns those permissions to all the team's members. Only administrators who can edit/manage the team can manage team permissions.

If you created a team whose task is to moderate the Message Boards, for example, you'd want to give the team all the permissions they'd need.

1. Go to *Site Administration* → *Content* → *Message Boards*.

2. Select *Home Category Permissions* from the *Options* icon ( ) in the top right of the screen.

3. Find the Team in the Role column and select the appropriate permissions.

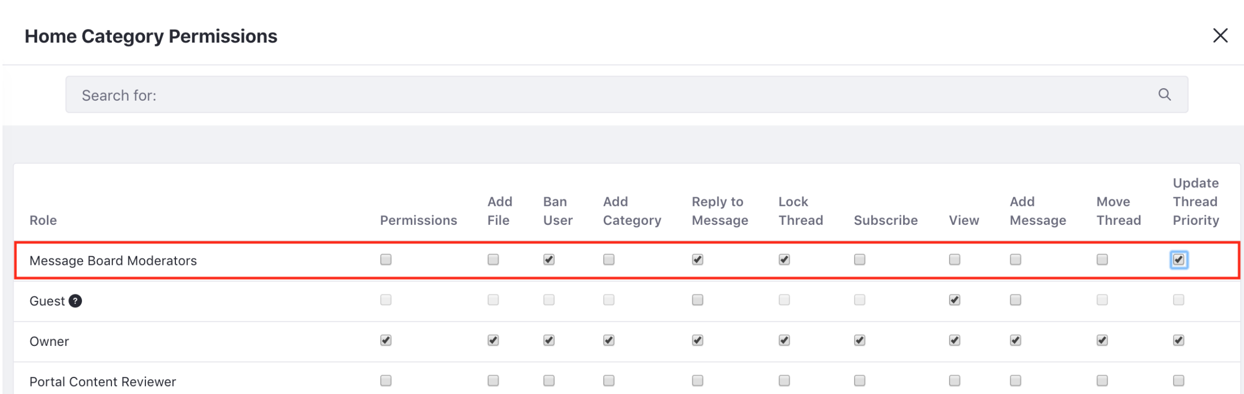

Figure 19.6: The Lunar Resort Message Board Moderators Site Team has unlimited permissions on the Message Boards application.

That's it! It's easy to give groups of Site Users permissions to perform their tasks.

These tutorials have introduced you to Liferay DXP Site management. You've learned how to use Liferay DXP to create multiple Sites with different membership types. You've also seen how easy it is to create and manage Sites and to create and manage pages within Sites. Next, you'll begin working with web content.

Chapter 20

Creating Web Content

Web Content Management (WCM) helps users who are not web developers publish content with a simple point and click interface, while enabling developers to create complex templates with dynamic elements. Once these templates have been deployed into Liferay DXP, your non-technical users can use them to manage complex content as easily as they would manage basic content.

It has these components:

Web Content Editor: A complete HTML editor for modifying fonts, adding color, inserting images, and much more.

Structure Editor: Define fields for structured content and more advanced designs.

Template Editor: Import template script files or create your own template to inform the system how to display the content within the fields determined by the structure.

Web Content Display: Place web content on pages in your Site.

Asset Publisher: Aggregate and display different types of content together in one view. This is covered in more detail in Publishing Assets.

Scheduler: Schedule when content is reviewed, displayed or removed. This is covered in more detail in Scheduling Web Content Publication.

Workflow Integration: Run your content through a review process. This is covered in more detail in the Workflow section.

Staging: Use a separate staging server or stage your content locally so you can keep your changes separate from the live site. This is covered in more detail in Using the Staging Environment.

These tools streamline the content creation process for end users and are also integrated with Liferay's services so advanced template developers can use them to query for data stored elsewhere on your website.

To demonstrate Liferay DXP's Web Content Management features, you'll create and manage content on Liferay for the ambitious (and fictitious) *Lunar Resort* project. The Lunar Resort project specializes in facilitating lunar vacations. It provides space shuttle transportation from the Earth to the Moon and back, offers the use of a state-of-the-art recreational facility enclosed by a large, transparent habitat dome, and even rents out lunar rovers.

CHAPTER 21

PUBLISHING BASIC WEB CONTENT

Web Content is one of many different kinds of assets, along with blog posts, wiki articles, message board posts, and other kinds of content. Like all of these assets, Liferay DXP handles Web Content using an asset framework that includes categories, tags, comments, ratings and more. Please see Publishing Content Dynamically for more information on Liferay's asset framework.

To start working with Web Content, publish some basic material using Web Content Management's WYSIWYG editor. Then you can cover the editor's features in greater depth and learn how to publish the content you create to a page.

21.1 Creating Web Content

To start you'll create and publish some simple content using the WYSIWYG editor to the home page of the Lunar Resort's web site.

1. Go to *Site Administration* and click the *Site Selector* button ().

   Content is created in whichever Site is selected, so always make sure that you're working on the right Site. Content can also be created as Global so that any Site can access it.

2. Select the *Lunar Resort* Site.

3. Open the *Content* section and click on *Web Content*.

   Here you can create web content and organize it into folders.

4. Click *Add* () → *Basic Web Content* to create a new web content article.

5. Type *Welcome to the Lunar Resort* in the *Title* field.

6. In the *Summary* field, give a short description of the Lunar Resort's facilities (be creative).

7. In the *Content* field, add the body of your web content article, which you'll dive into next.

8. Finally, leave the *Searchable* switch enabled and click *Publish*.

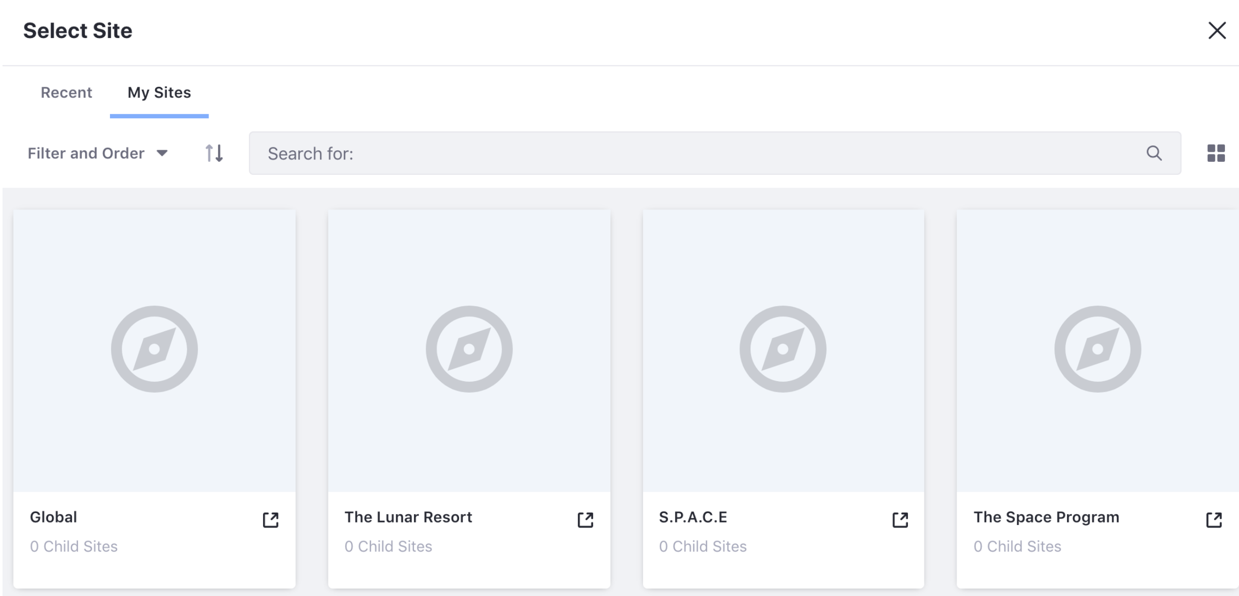

Figure 21.1: You can choose where to create content by navigating to the Site Administration menu and selecting your Site and page scope.

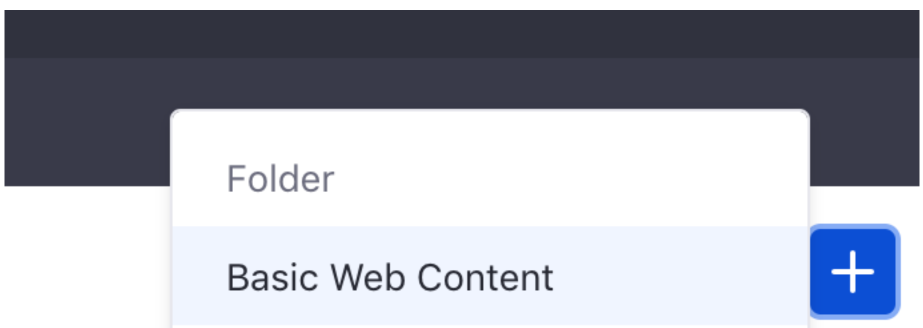

Figure 21.2: By default, *Basic Web Content* is the only article type available. The next tutorial covers how to create new types.

> Disabling search for your article removes it from end users' search results. Administrators can still search for it from *Site Administration* → *Web Content*, and the article can still be added to pages.

That's all it takes to create Web Content, but there's much more under the hood of the web content editor that you'll learn about next.

21.2 Using the Web Content Editor

In the previous tutorial you created a simple web content article, but the Web Content editor can do much more than make plain text articles.

Basic Editor Functions

To explore these options, go back to the article you just created and make it better:

1. On the *Web Content* page in *Site Administration*, click on the title of the article.
2. Highlight the text you entered in the *Summary* field.

   A number of controls appear. These let you style the text or provide a link.
3. Click on the arrows where it says *normal* to open the *Styles* dropdown and select the *Heading 1* style.

Next, you'll add an image to the article. Whenever you place your cursor in the *Content* area, the *Add* icon (+) appears. If you click on it, controls for inserting an image, video, table, or horizontal line ( ) appear.

1. Click *Add* (+).
2. Select the icon that depicts a mountain silhouette to insert the image.
3. In the image file selector, select an image to add to the article.

   Select an image from your computer or from the Site's Documents and Media repository. If you select one from your Documents and Media repository, you can access the image editor to make changes specifically for your article.

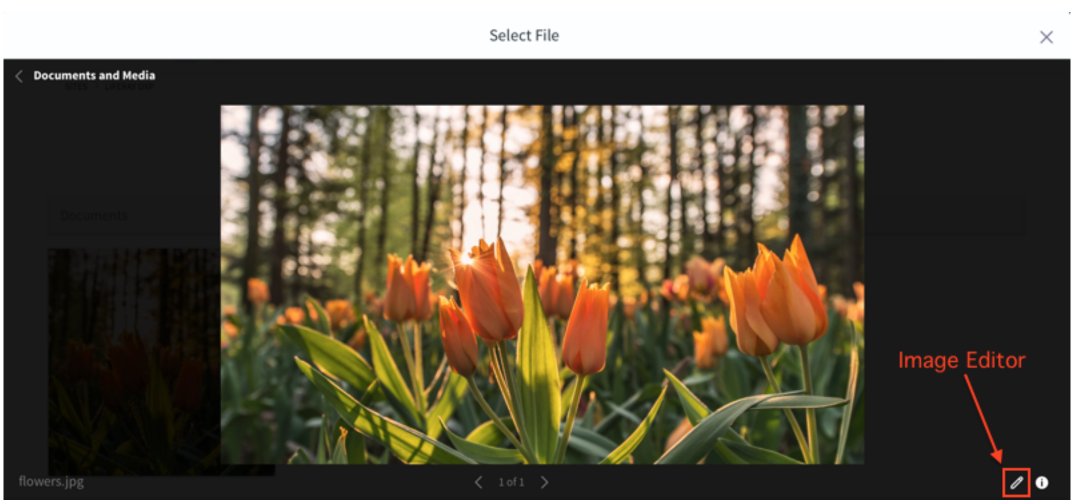

Figure 21.3: You can access the image editor through the item selector window.

After adding an image to the web content article, click it to bring up controls ( ) for formatting it. You can also make it a link.

The same way you inserted an image in to the article, you can also insert a table. Click the table to access edit controls, which let you designate the first row and/or column as table headers, and also enable you to add rows, columns, and cells.

In addition to images and tables, you can insert a horizontal line as a separator between between sections. You can also add video by providing a URL.

Editing the Article Source

If you need to work directly with the HTML, you can switch to source view.

1. With your cursor in the *Content* field, select the *Source* icon ( ) to switch.

2. Click the regular mode icon ( ) to go back once you're done editing HTML.

The HTML editor highlights syntax, and you can switch between a dark and light theme by choosing the moon and sun icons.

In HTML mode, click on the *Fullscreen* icon ( ) to access a dual pane view that shows your HTML code on the left and a preview pane on the right. You can arrange the HTML and preview panes horizontally or vertically.

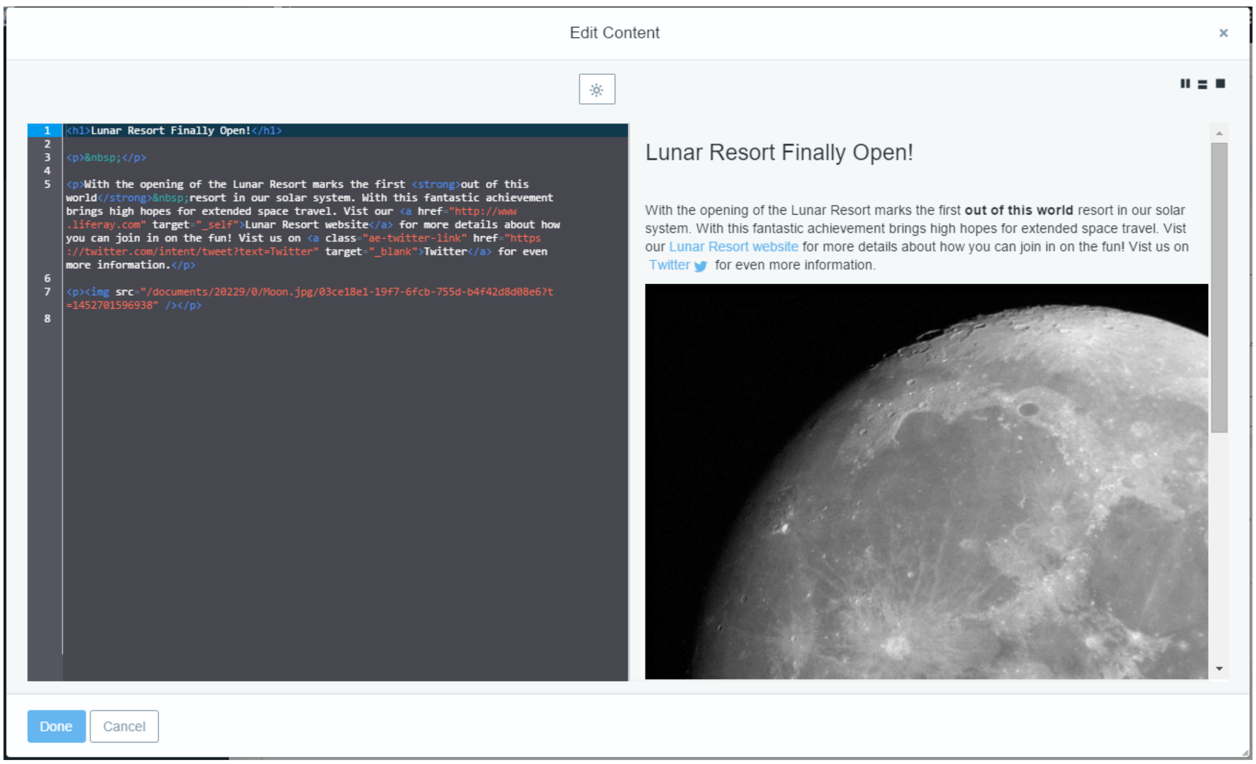

Figure 21.4: You can view how your HTML would render by using the preview pane.

You can exit the enlarged editor by clicking the *Done* button at the bottom of the screen.

1. Add a few short sentences announcing the grand opening of the Lunar Resort.

2. Click *Save as Draft*.

   Be sure to save your content frequently, because it is not auto-saved.

The content can be localized in whatever language you want. You'll learn more about localizing your content later.

Web Content Options

The bottom menu of the New Web Content form provides options for customizing your web content.

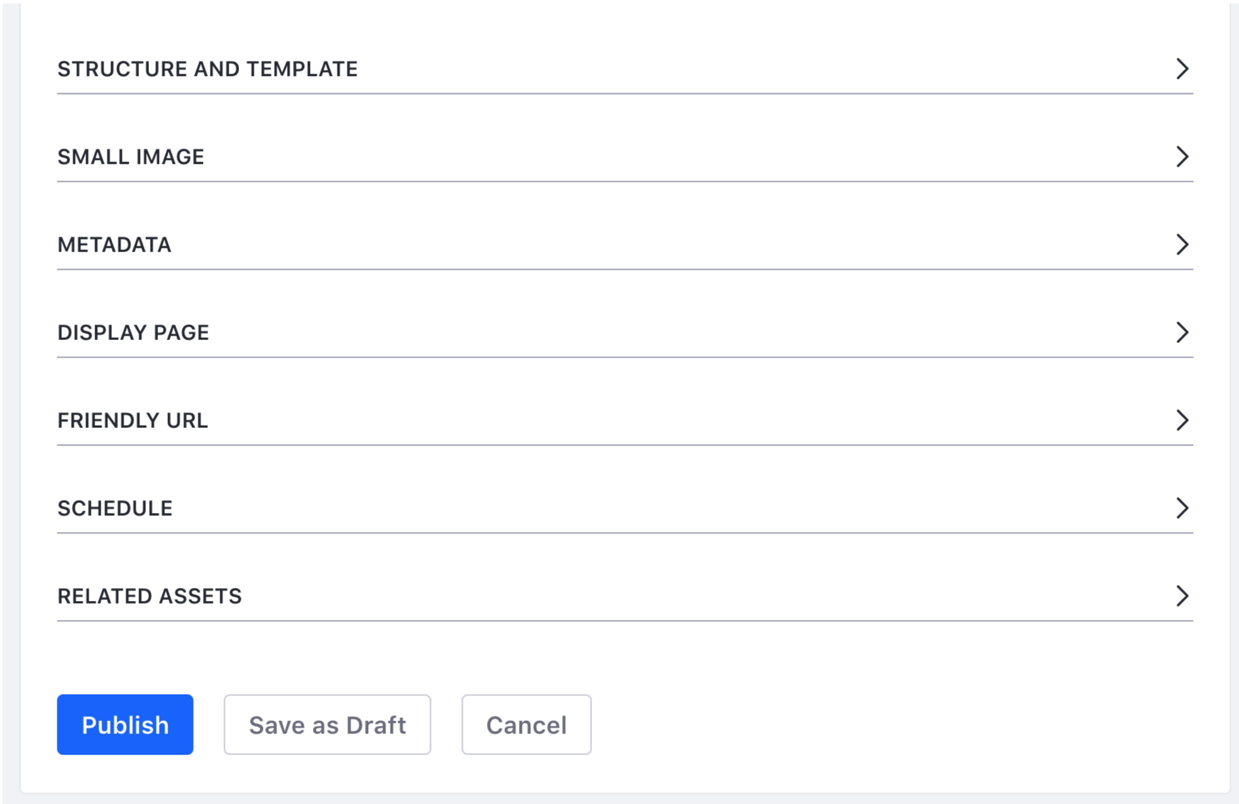

Figure 21.5: New web content can be customized in various ways using the menu located below the editor.

Structure and Template: Customize the web content article's structure and template. To learn more about web content structures and templates, see Designing Uniform Content.

Small Image: Set the image that is used for the web content article's previews.

Metadata: Organize web content articles by selecting tags, categories, and priority. To learn more about tags and categories, see Organizing Content with Tags and Categories.

Display Page: Select a Display Page to enhance the styling and formatting of your web content.

For example, if you had a news site with different sections—Sports, Technology, Culture—you could create a display page for each section with unique banners, formatting, embedded widgets, or other features. By selecting a display page, you would ensure that content appears on the page with the appropriate features. You'll work through an example of creating a display page in the Display Pages for Web Content tutorial.

Friendly URL: Set the friendly URL where the article can be viewed alone. If a specific display page is set, the URL links to it.

Schedule: Customize the date and time your content publishes and/or expires. To learn more about scheduling content, see Scheduling Web Content Publication.

Related Assets: Determine relationships between the web content article and other assets, even if they don't share any tags and aren't in the same category. You can connect your content to any asset that implements the Related Assets feature. To learn more about defining content relationships and publishing links to those related assets, see Defining Content Relationships.

Figure 21.6: This blog entry has links to two Related Assets: an article and a message board thread.

Permissions: Customize who has access to the content. By default, content is viewable by Anyone (Guest Role). You can limit viewable permissions by selecting any Role from the drop-down or in the list. You can customize permissions in more detail by selecting the *More Options* link below the drop down button.

While you can set permissions here, they are ignored unless you activate Web Content Article permissions in your System Configuration:

1. Go to *Control Panel* → *Configuration* → *System Settings*.
2. Search or browse for *Web Content (Default Settings for All Instances)*.
3. Check the box labeled *Article View Permissions Check Enabled*.
4. Click *Save*.

Once it is activated, any permissions you set in the article's configuration are checked before displaying the article.

21.3 Publishing Web Content

In the previous sections, you created and edited an article. Now it's time to publish it.

1. Go to the *Welcome* page of the Lunar Resort Site.
2. Select the *Add* button (+) from the top Control Menu and select the *Widget* tab.
3. Find *Web Content Display* and add it to the page.

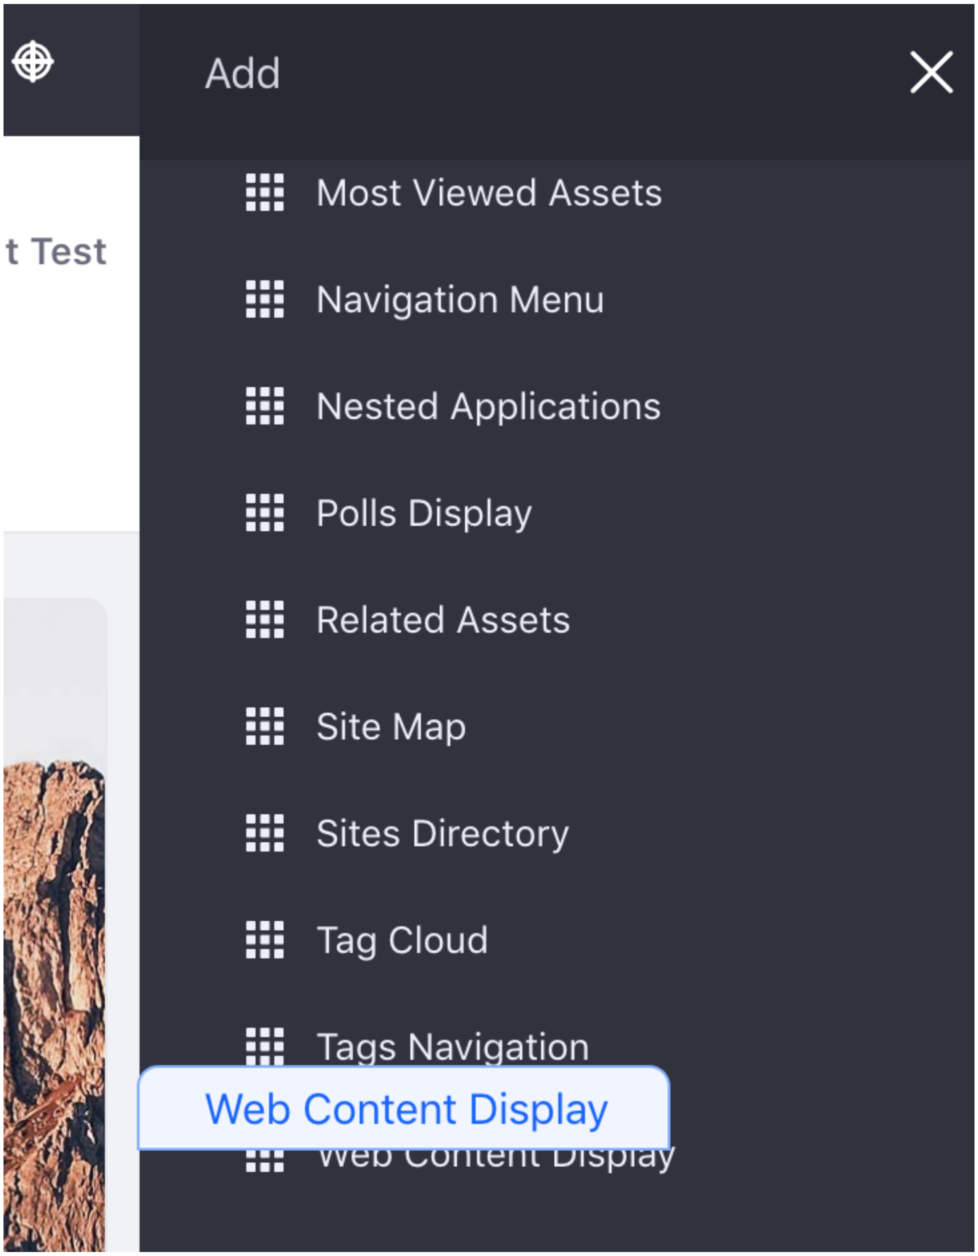

You can drag a widget to the position on the page where you want your content to appear. You can have as many Web Content Display widgets on a page as you need, which gives you the power to lay out your content exactly the way you want it.

Now select the content to display:

1. Click *Select Web Content to make it visible* in the bottom of the widget.

2. Click *Select* under *Web Content*.

3. Click on the article that you want to display

   If your content does not immediately appear in the list, you can search for the content by title, description, user name, or Site (click the drop-down arrow to see all the options).

Selecting a web content article displays the Web Content Display's configuration page, where you can choose the User Tools and Content Metadata to be published in the widget. These two entities have the following options to choose from, by default:

- **User Tools**
  - *Translations*
  - *Print*
- **Content Metadata**
  - *Related Assets*
  - *Ratings*
  - *Comments*
  - *Comment Ratings*

If you have enabled OpenOffice/LibreOffice integration, you can also enable document conversion for your content. Then users can download your content in their format of choice. To enable OpenOffice/LibreOffice integration, go to *Control Panel → Configuration → System Settings → Connectors* and check the *Server Enabled* box. Back in the Web Content Display's configuration page, conversion options are available under the *User Tools* list.

Translations: Shows the available locales for your content. If you're working on the page for a particular language, you can select the translation of your content that goes with your locale.

Print: Opens the content in a separate browser window with just the content—no navigation or other widgets.

Ratings: Shows one of two ratings interfaces Liferay DXP has: five stars or thumbs up and thumbs down. Rating types are configured in *Site Menu → Configuration → Site Settings*.

Comments: Creates a discussion forum attached to your content.

By default, guests cannot leave comments on web content. If you want to allow guests to comment on your web content article,

1. Navigate to the Control Panel → *Users → Roles*
2. Select *Guest → Define Permissions*.
3. From the left menu, select *Site Administration → Content → Web Content*.
4. Navigate down to the Web Content Article heading and select the *Add Discussion* checkbox. Click *Save*.

Guests can now post comments on your web content article!

You may decide you want one, some, or none of these features, which is why they're all implemented as simple selector buttons to be enabled or disabled at need. Once you've selected the features you want to include in your Web Content Display widget, click *Save* and close the configuration window.

Editing Published Content

If you must edit published content, you can do it directly from the Web Content Display app or from Site Administration. To edit it from the Web Content Display app,

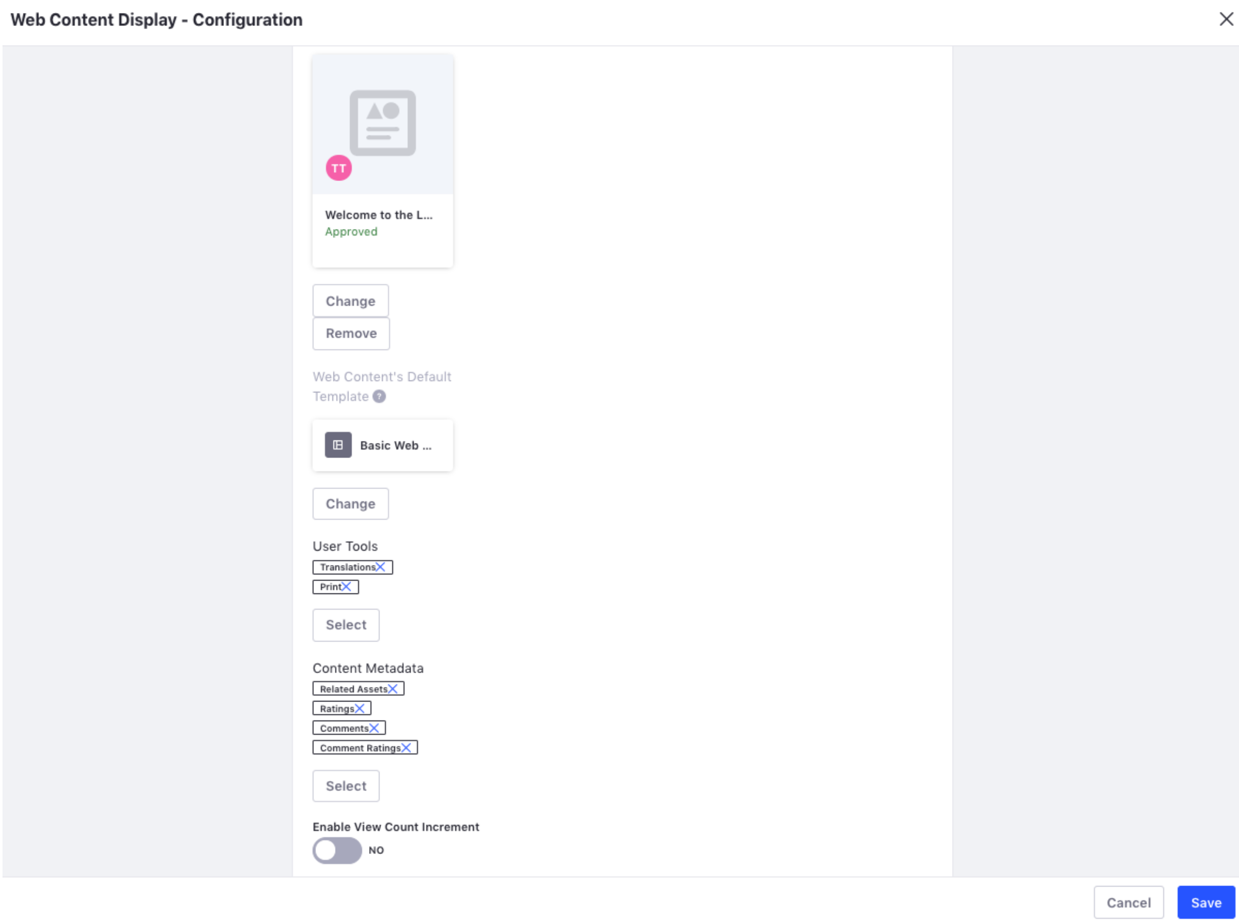

Figure 21.7: Publishing web content is a snap. At a minimum, you only have to select the content you wish to publish. You can also enable lots of optional features to let your users interact with your content.

1. Select the *Options* button ( ) from the widget's top panel.
2. Select *Edit Web Content* to launch the editor. Select *Edit Template* to launch the template editor for the web content article's template.

If you edit the article from Site Administration, you can also view the article's history and use the diff tool to compare versions.

1. Go to *Content* → *Web Content* from the Menu.
2. Next to the article, click *Actions* icon ( ) and select *View History*. This shows you

This shows you all the article's versions and modified/display dates. The diff tool compares these versions and highlights the differences between them.

1. Click *Actions* next to a version of the article you'd like to compare.
2. Select *Compare to...*.

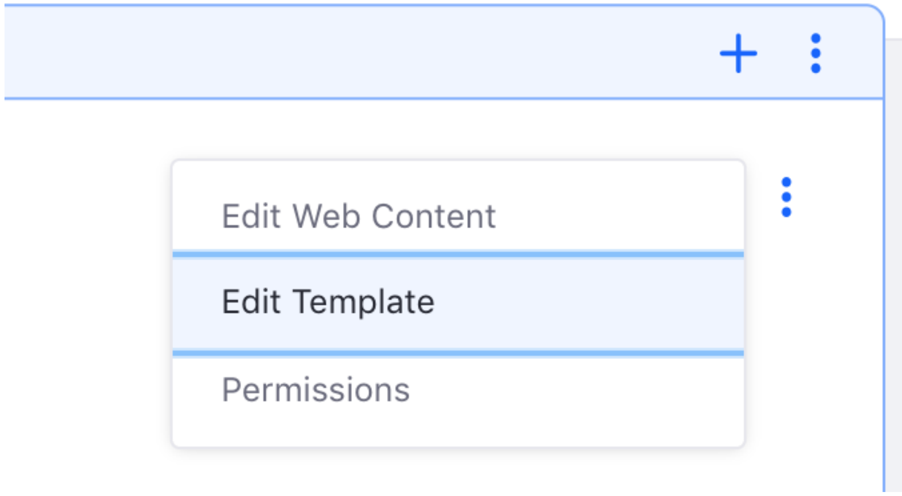

Figure 21.8: You can select and edit an article, or edit its template directly from the Web Content Display app.

3. Select the version with which to compare it.

   The tool provides color coded highlighting to emphasize additions and deletions between the two articles.

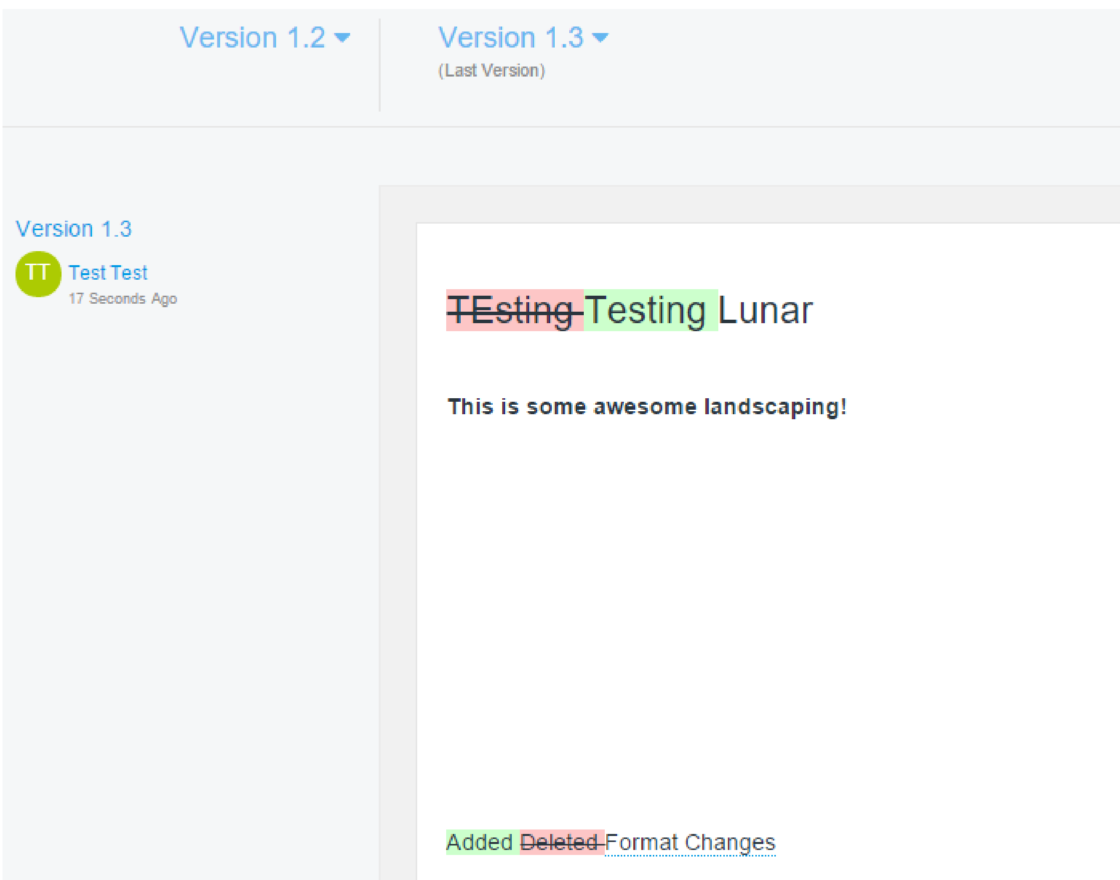

Whenever you publish updates to a web content article that's already being displayed, the content is immediately updated unless you have a workflow enabled (See Workflow for details).

21.4 Other Content Options

Here are some options and tools that you can use to enhance your content and user experience.

Localizing Content

When you create a new web content article, you can choose a default language. First, you must change the system configuration to enable the option to change the default language.

1. Go to the *Control Panel* → *System Settings*.
2. Locate *Web Content* → *Administration* by scrolling or using the search bar.
3. Check the box labeled *Changeable Default Language*.
4. Click *Save*.

You must now add translations for any languages you need. Adding translations works like this:

1. Open a web content article.
2. Click the flag icon with a country code on it next to any localizable web content field.
3. Select a language from the list.

When you select a language, all fields in the article switch to the new language. To create the new translation, fill in the fields in the selected language and publish the article.

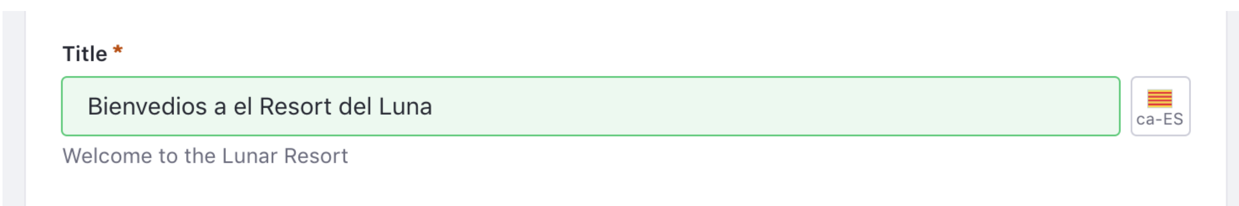

Figure 21.9: Adding a translation to an article works like adding the default translation.

Note: To view localizable fields in a given language, you must have your Portal set to that language. This includes friendly URLs for the web content as well. When you navigate to the localized friendly URL (e.g. `http://localhost:8080/web/guest/-/espanol`), the web content is always displayed in the current language. You can change the language with the Language Selector app.

You can modify the language translation list by inserting `locales.enabled=` followed by your preferred languages in your `portal-ext.properties` file. For example, `locales.enabled=ar_SA,nl_NL,hi_IN` offers *Arabic (Saudi Arabia)*, *Dutch (Netherlands)*, and *Hindi (India)*.

Warning: If you switch your Site's default language (e.g., via friendly URL), but do not have the necessary translations for localizable fields, your Site's language values are used from the old default language. Therefore, you should change the default language of your Site *only* when you have translated values for all localizable entities.

When you create a new web content structure, each field you create has a *Localizable* checkbox displayed next to it. This enables you to control what can and can't be changed in the translation process. For example, if you don't want images or content titles to be changed when the content is translated, you can make sure those fields aren't llocalizable. When you follow the steps above to localize content, only fields within the structure that had the *Localizable* box checked appear within the translation window.

Xuggler for Embedding Video

Xuggler is a tool which generates video previews and makes it possible to embed videos from your Documents and Media library in web content and elsewhere on the site. To enable Xuggler,

1. Navigate to the *Control Panel*.
2. Click on *Configuration* → *Server Administration* → *External Services*.
3. Scroll to the bottom and click *Install* in the *Xuggler* section.

   This downloads the necessary libraries and prompts you to restart the server to enable Xuggler.
4. After the server restarts, you can enable Xuggler from the same page.

Once Xuggler has been installed and enabled, you can embed a video or audio file in a web content article the same way you added images previously.

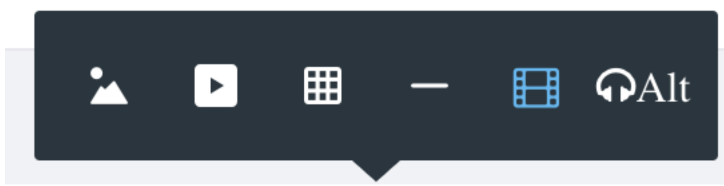

Figure 21.10: If you've installed and enabled Xuggler from the *Server Administration* → *External Tools* section of the Control Panel, you can add audio and video to your web content!

XML Format Downloads

Tools like the Resource Importer and Site Initiators can be deployed to build a site almost instantly. Before you can use them to import Web Content, however, you first need to have the content exported individually in XML format. To export the content,

1. Go to *Site Administration* → *Content* → *Web Content*.
2. Start editing the article you want to download.
3. Click the *Options* icon ( ) in the top right of the page and select *View Source*.

This displays the raw XML source of the article. You can copy this content to save into an XML file locally.

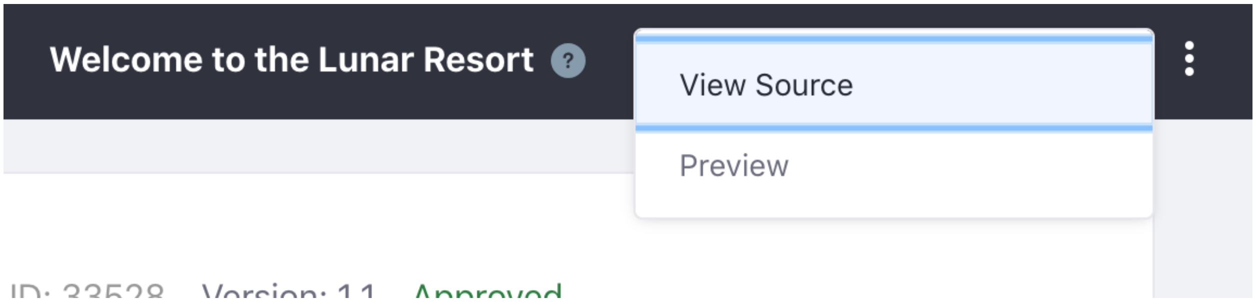

Figure 21.11: The *View Source* button is available from the *Options* button.

Subscribing to Content

An administrator or web content writer can subscribe to an article or folder to follow changes being made to it.

1. Go to *Content* → *Web Content* for your Site.

2. Click *Options* () → *Subscribe* next to the article or folder you want to follow.

Anytime an asset that you follow is modified, you receive an email notifying you of the change.

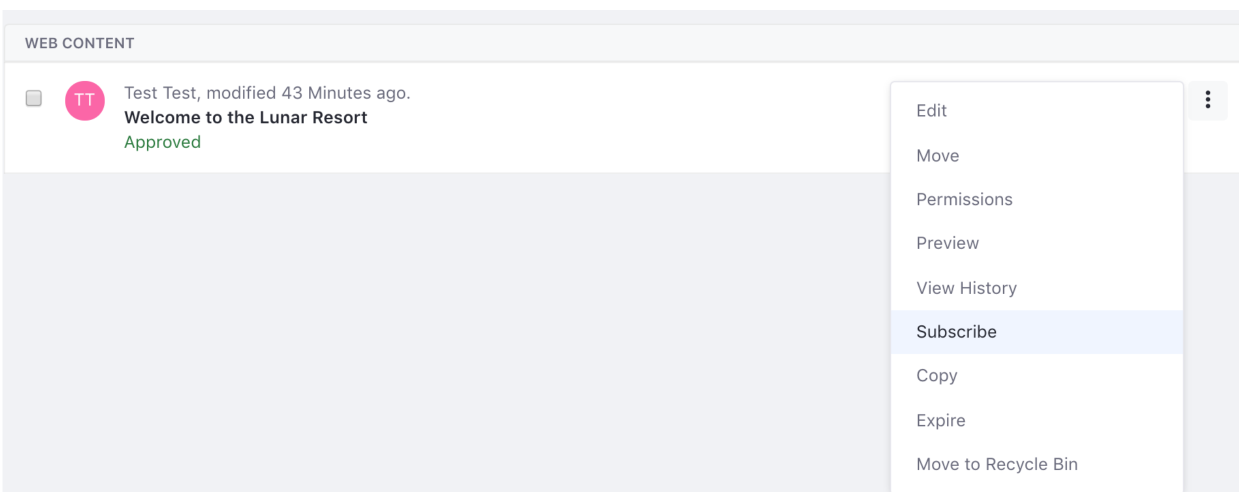

Figure 21.12: Click the Subscribe icon in the web content entity's *Options* menu to begin receiving web content notifications.

That's pretty much all there is to basic content creation. Whole sites have been created this way. But if you want to take advantage of the full power of Liferay DXP's WCM, you'll want to use structures and templates or Fragments. You'll cover these topics next.

Organizing Structure Names

By default, when you select a structure to add a new Web Content article, the structures are ordered by their IDs, not their names. This can be confusing, but—never fear—there's a configuration property to sort them alphabetically.

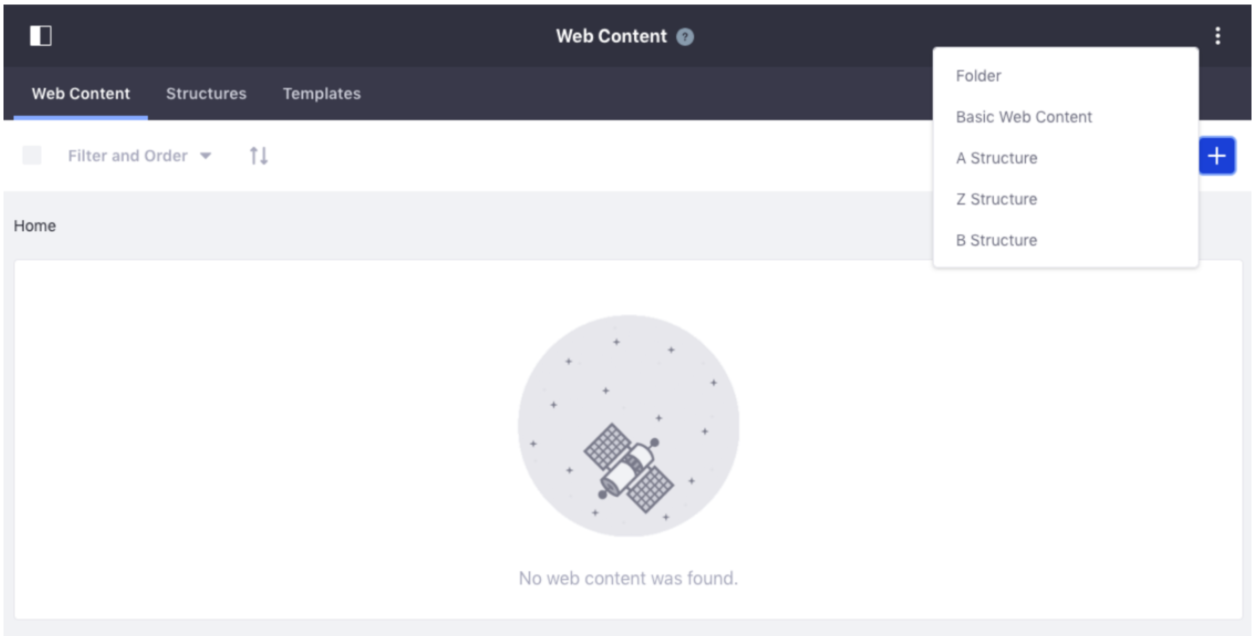

Figure 21.13: The default ordering for Web Content Structures can yield confusing results.

To enable this property for Site Administration,

1. Go to *Configuration* → *System Settings* → *Web Content* → *Web Content Administration*.
2. Check the box labeled *Journal Browse by Structures Sorted by Name*.

With 7.0 Fix Pack 2 or greater, you can also set this property for the Web Content Display widget. To enable this property for the Web Content Display,

1. Go to *Configuration* → *System Settings* → *Web Content* → *Web Content Display*.
2. Check the box labeled *Sort Structures by Name*.

After this option is checked, the structures are sorted alphabetically. Note that enabling this property can degrade performance with large structure libraries.

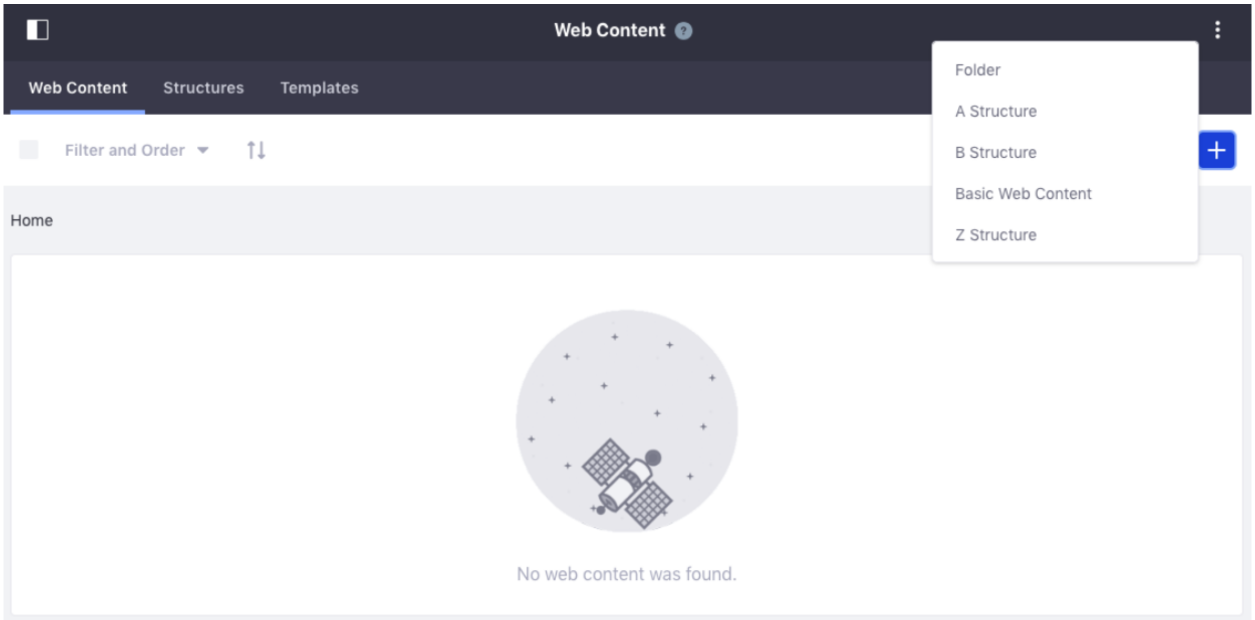

Figure 21.14: Web Content Administration will now display structures in alphabetical order.

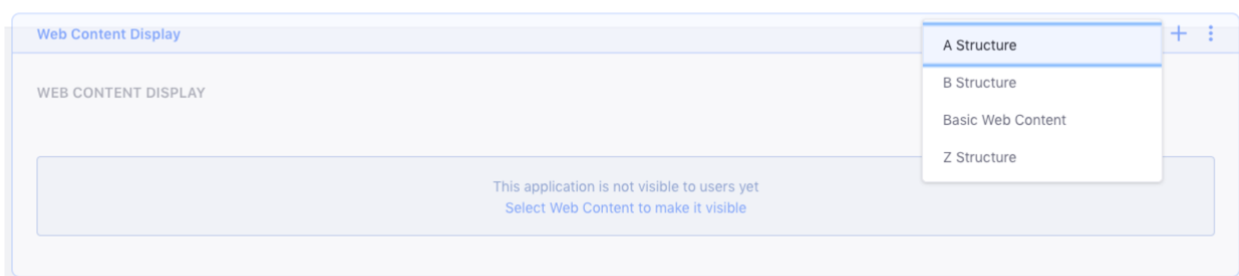

Figure 21.15: The Web Content Display widget will now display structures in alphabetical order.

Chapter 22

Designing Structured Content

If you've ever launched a web site, you know that as it grows, you can experience growing pains. This is the case especially if you've given lots of people access to the site to make whatever changes they need to make. Without preset limitations, users can display content in any order and in any manner they desire (think huge, flashing letters in a font nobody can read). Content can get stale, especially if those responsible for it don't maintain it like they should. And sometimes, content is published that should never have seen the light of day.

!PVideo Thumbnail

Thankfully, Web Content Management helps you handle all of these situations. You can use *Structures* to define which fields are available to users when they create content. These can be coupled with *Templates* that define how to display that content. Content won't get stale, because you can take advantage of the Scheduling feature to determine when content is displayed and when it's removed. Additionally, you can configure Liferay DXP's built-in Workflow system to set up a review and publishing process so only what you want winds up on the live site. This gives you what you need to run everything from a simple, one-page web site to an enormous, content-rich site.

All of this starts with structures.

!VVideo Tutorial

CHAPTER 23

CREATING STRUCTURED WEB CONTENT

Structures are the foundation for web content. They determine which fields are available to users as they create new items for display. Structures not only improve manageability for the administrator, they also make it much easier for users to add content quickly.

For example, say you're managing an online news magazine. All your articles must contain the same types of information: a title, a subtitle, an author, and one or more pages of text and images that comprise the body of the article. With only basic content creation, you'd have no way to make sure your users entered a title, subtitle, and author. You might also get articles that don't match the look and feel of your site. If titles are supposed to be navy blue but they come in from your writers manually set to light blue, you must spend time reformatting them before they are published.

Structures enforce a format for your content so your writers know exactly what a complete article needs. Using structures, you can provide a form for your users which spells out exactly what is required and can be formatted automatically using a template.

You create a structure by adding form controls such as text fields, text boxes, HTML text areas, check boxes, select boxes and multi-selection lists. You can add specialized application fields such as an Image Uploader or Documents and Media right onto the structure. Positioning elements is accomplished by drag-and-drop, making it easy for you to prototype different orders for your input fields. Additionally, elements can be grouped together into blocks which can then be repeatable. Template writers can then write a template which loops through these blocks and presents your content in innovative ways, such as in sliding navigation bars, content that scrolls with the user, and more.

Next you'll learn how you can create and edit structures through the Manage Structures interface.

23.1 Editing Structures

To start, go to the *Structures* page.

1. From *Site Administration* go to *Content* → *Web Content*.

2. Open the *Structures* tab.

This page shows you all the web content structures in this Site. You can add new web content structures, edit existing ones, manage the templates associated with a structure, edit the permissions of a structure, and copy or delete structures.

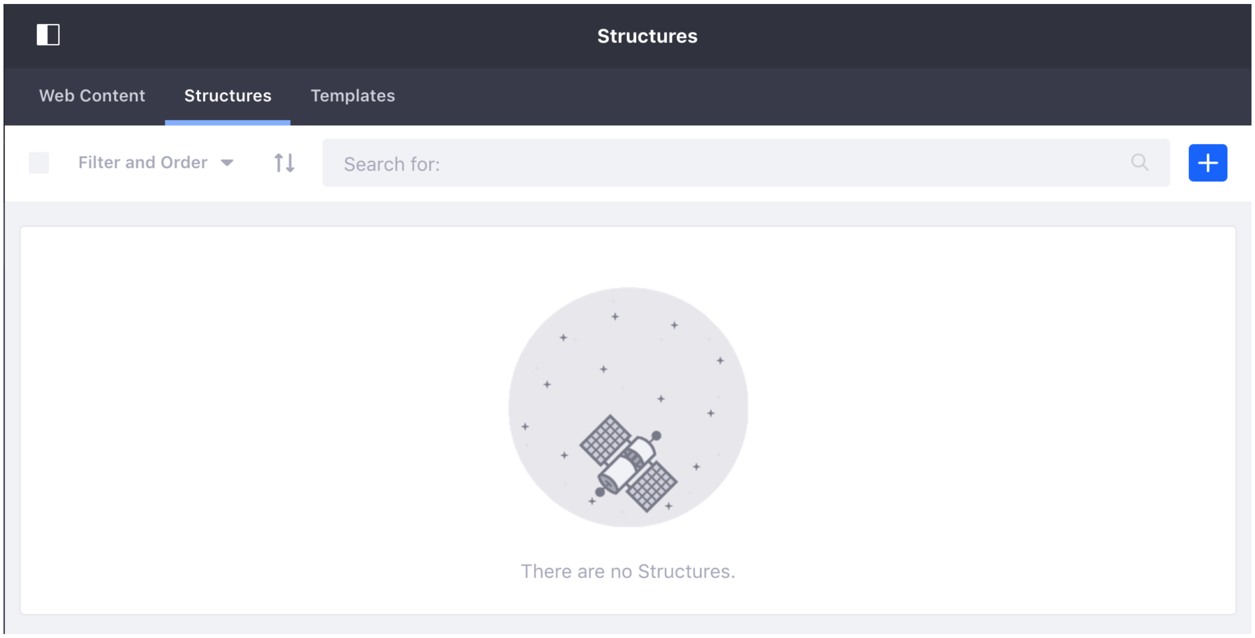

Figure 23.1: Structures are not pre-installed. You have to make your own.

Note: When you copy a structure, Liferay DXP generates a unique ID for the copied structure, but every other attribute of the copied structure, including the name, is the same as that of the original. When you copy web content structure, enter a new name for it to avoid confusing it with the original. During the copy process, you're prompted to choose whether to copy any detail templates or list templates associated with the structure. For information on detail templates and list templates, please refer to Dynamic Data Lists.

Basic Web Content which you used in previous exercises lives at the *Global* scope so that it is available to all Sites. This structure and template are used automatically if a custom structure and template are not added.

Structure Fields

Now, create a new Structure:

1. Click *Add* ().
2. Give your Structure a name.

Structures are essentially a set of fields organized in a certain way. The interface on this page provides an easy way to add and organize whatever fields you need. Each element that you add has three icon options that you can click:

Settings: ( ) Changes the name and label and set other information about the field, like whether or not it is required.

Delete: ( ) Removes the field from the structure.

Duplicate: ( ) Duplicates the field and all its settings and iterates the *Name* to avoid conflicts.

Web content structures can inherit characteristics from other structures. A child structure inherits all the parent's fields and settings. You can use this to make a similar structure to one that already exists. For example, if you have *Sports Article* and you want to create *In-depth Sports Article*, set *Sports Article* as the parent and the *In-dept Sports Article* inherits all its fields, letting you add new ones for more in-depth information.

Note: Due to import/export operations it's possible to have both a global and a Site-scoped structure with the same `structureKey`. If this happens, the Site-scoped structure takes precedence, and you can't access the global structure from that Site.

You can also manually customize a structure's XML in *Source* mode. By default the *View* mode is selected, but you can click the *Source* tab to switch. This method is for more experienced developers.

Take a moment to add, delete, and rearrange different elements.

The following fields can be in structures:

Boolean: Adds a checkbox onto your structure, which stores either `true` (checked) or `false` (unchecked). Template developers can use this as a display rule.

Date: Adds a pre-formatted text field that displays a date picker to assist in selecting the desired data. The format for the date is governed by the current locale.

Decimal: Similar to *Number*, except that it requires a decimal point (.) be present.

Documents and Media: Adds an existing uploaded document to attach to the structure. Can also upload documents into the Document Library.

Geolocation: Adds a map that displays a configured location. The geolocation system can work in two ways: letting the system know your current location (especially useful on mobile devices) and giving the user directions to a another place.

HTML: An area that uses a WYSIWYG editor to enhance the content.

Image: Adds the browse image application into your structure. You can select an image from the Documents and Media library or upload an image from your computer's storage. If uploading an image from your personal computer to the web content article, it is only available for that article.

Integer: Similar to *Number*, except that it constrains user input to whole numbers.

Link to Page: Inserts a link to another page in the same site.

Number: Presents a text box that only accepts numbers as inputs, but puts no constraints on the kind of number entered.

Radio: Presents the user with a list of options to choose from using radio button inputs.

Select: Presents a selection of options for the user to choose from using a combo box. Can be configured to allow multiple selections, unlike *Radio*.

Separator: Adds a horizontal line between fields.

Text: Used for items such as titles and headings.

Text Box: Used for the body of your content or long descriptions.

These fields provide all you need to model any information type you would want to use as web content. Liferay customers have used structures to model everything from articles, to video metadata, to wildlife databases.

23.2 Configuring Structure Fields

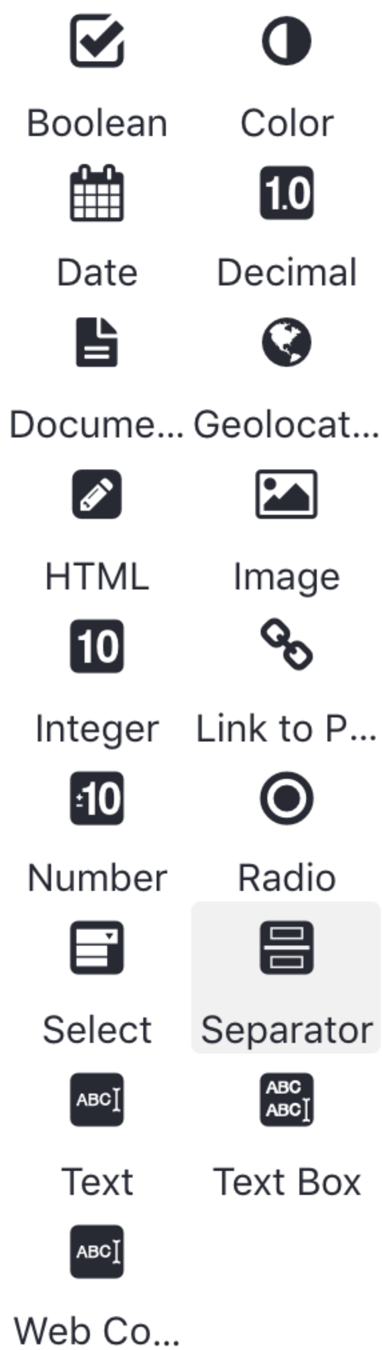

Figure 23.2: The structure editor gives you many options to customize your Web Content.

There are many options available for configuring each structure field. Some of them relate to how the fields are displayed or how users interact with them, but probably the most important field configuration is the **Name**. When you create a new field, it has a random name generated that looks like `TextField4882`. In most cases, you should change this to something that is more memorable and more descriptive. When it comes time to create the template, you don't want to be trying to remember if `TextField4882` was the field for entering an applicant's name or annual salary.

Practice this now.

1. In your structure, add an *HTML* element.

2. Hover over the field and select the *Configuration* icon ( ).

3. Change the *Field Label* value to *Instructions* and the *Name* value to *steps*. Now your template writer has a variable by which he or she can refer to this field.

Here's a list of all the configurable settings available for a structure's fields:

Type: Lists the type of field placed in the definition. This is not editable but is available to reference from a template.

Field Label: Sets the text that can be displayed with the field. This is the human-readable text that the user sees.

Show Label: Select *Yes* to display the Field Label.

Required: Select *Yes* to mark the field required. If a field is required, users must enter a value for it in order to submit content using this structure.

Name: The name of the field internally, automatically generated. Since this is the variable name that you can read the data from in a template, you should enter a descriptive name.

Predefined Value: When a user creates a new web content article based on a structure that has predefined values for various fields, the predefined values appear in the form as defaults for those fields.

Tip: Each field can have a small help icon, with a tooltip attached that displays helpful information. If you want to provide text for the tooltip, you may enter it here.

Indexable: Select *Yes* to permit your field to be indexed for search.

Localizable: Select *Yes* to permit localization for this field.

Repeatable: Select *Yes* to make your field repeatable. Users can then add as many copies of this field as they need. For example, if you're creating a structure for articles, you might want a repeatable Author field in case you have multiple authors for a particular article.

Multiple: Select *Yes* to enable a multi-selection list (only available for the Select field).

Options: Changes the options available for selection. You can add and remove options as well as edit each individual option's display name and value (only available for Radio and Select fields).

Style: Changes the line separator's style (only available for Separator).

Structure Default Values

You can define Structure Default Values for repeatable values in content created from that structure. They can also set defaults for Liferay's standard asset fields (like tags, categories, and related assets) and the content of the structure fields, while also setting a default template for displaying the structure data.

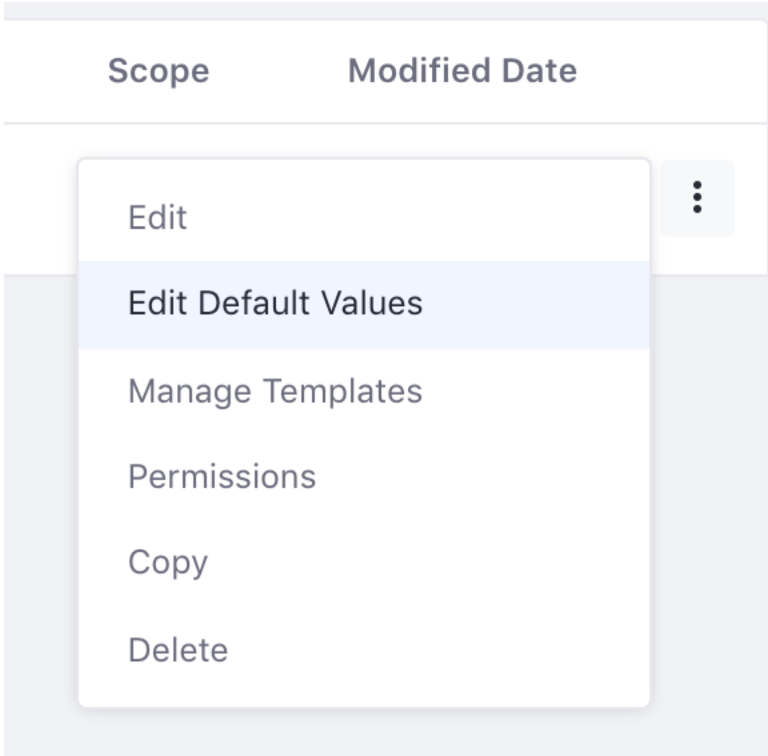

Figure 23.3: You can edit default values via the *Actions* button of the Manage Structures interface.

Returning to the newspaper scenario again, suppose you want all sports articles to have the same display page (sports page), the same categories, or the same set of tags. Instead of adding them for each article or wondering if your users are adding them to every web content article, you can add these characteristics once for every sports article by creating default values for the structure. Creating default values is not part of creating a new structure, so make sure you have an existing structure.

To edit a structure's default values:

1. Go to *Site Administration* → *Content* → *Web Content* and click on the *Structures* tab to see the structures list.

2. Find the *Actions* button (⋮) for the desired structure and select *Edit Default Values* from the menu to view a window like the one below.

   This form manages the structure settings. It duplicates the function of the *Predefined Value* field setting (see above), but is much more convenient for setting or editing a large number of defaults at once.

Every new web content you create with this structure is preloaded with the data you inserted. Next, you'll learn about assigning permissions.

23.3 Structure Settings

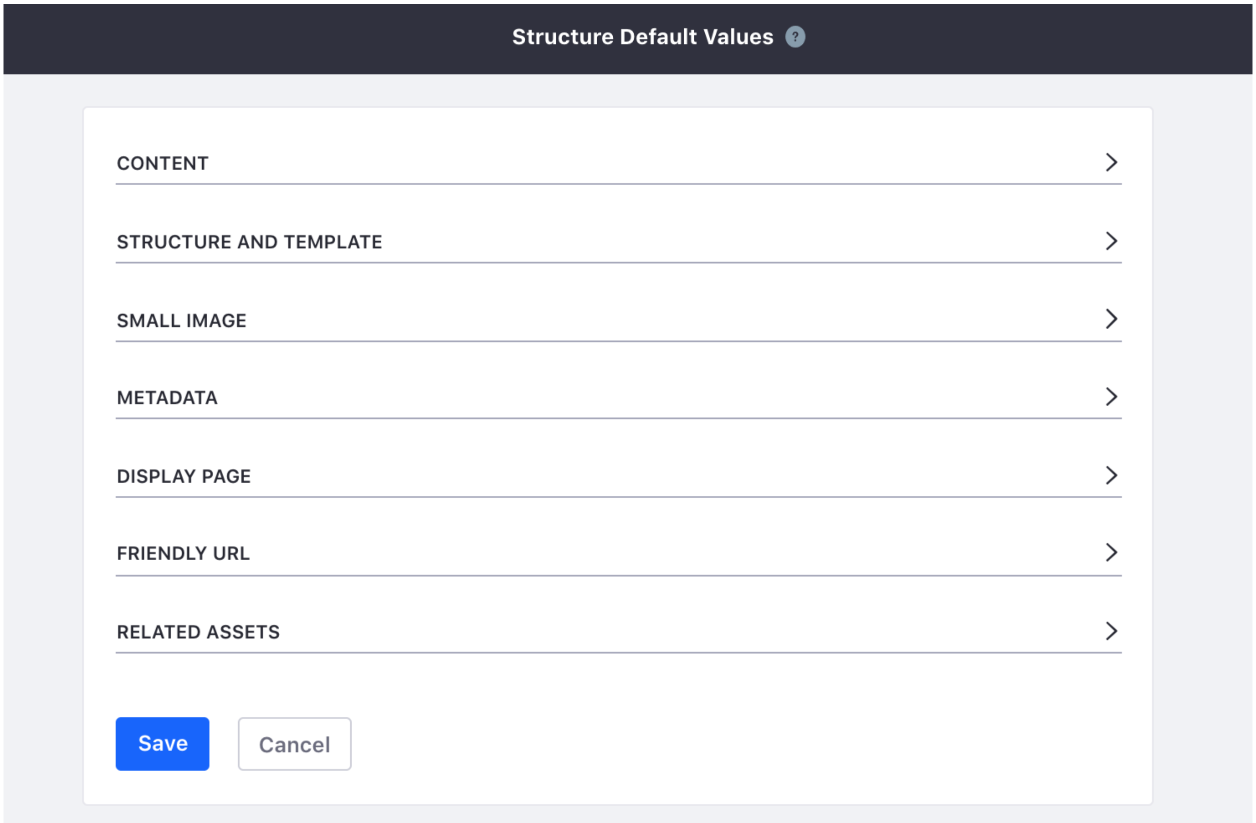

Figure 23.4: You can define values for your structure fields and the standard asset metadata fields.

After you have created a structure and configured its fields, there is additional configuration to consider including permissions and remote access to managing structures.

Assigning Permissions

Permissions on structures can be set just like any other permission. Most users should not be able to edit structures. Structures are coupled with templates, which require some web development knowledge to create. This is why only trusted developers should be able to create structures and templates. Users, of course, should be able to view structures. The *View* permission enables them to make use of the structures to create content.

The best practice for structure permissions is to grant or deny them based on Roles.

WebDAV URL

The WebDAV URL feature is available for web content structures and templates so users can upload and organize resources from both a web interface and the file explorer of their desktop operating system. With the WebDAV URL, site administrators can add, browse, edit, and delete structures and templates on a remote server. After you complete your structure, you can access the WebDAV

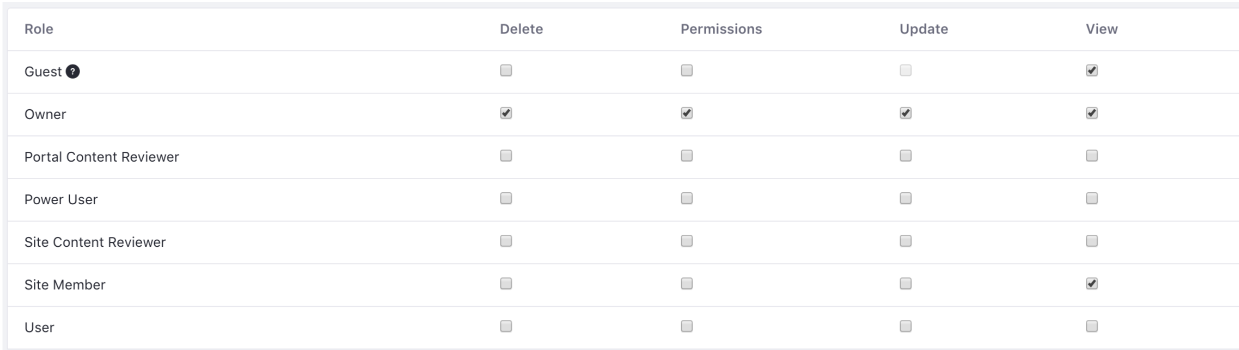

Figure 23.5: You're able to assign structure permissions via the *Actions* button.

URL by re-opening the structure or template and clicking the *Details* section. If you'd like the see WebDAV in action, see WebDAV Access.

Note: Some operating systems require a WebDAV server to be class level 2 (to support file locking) before allowing files to be read or written. The Documents and Media library uses a class level 2 WebDAV server but Web Content structures and templates do not. This means that Liferay DXP's Document and Media library supports WebDAV file locking but Web Content structures and templates do not. However, on operating systems which require WebDAV servers to be class level 2, it's possible to avoid the restriction by using third-party WebDAV clients (e.g., Cyberduck).

Now that you understand how structures work, you're ready to understand the other half of Liferay DXP's web content management system: templates.

CHAPTER 24

DESIGNING WEB CONTENT WITH TEMPLATES

Developers create templates to display the elements of the structure in the markup they want. In essence, templates are scripts that tell Liferay DXP how to display content in the structure. Changes to the structure require corresponding changes to the template, because new or deleted fields produce errors on the page. A structure cannot be displayed without a template, and a template cannot exist without a structure.

Liferay DXP supports creating web content templates in Freemarker (FTL), Velocity (VM), and XSL format. Freemarker is the preferred, recommended language, and what you'll use in the next example.

24.1 Adding Templates with Structures

To better understand templates, now you'll create a structure and an associated template. First create the structure:

1. Go to *Content* → *Web Content* from Site Administration page and open the *Structures* tab.

2. Click the *Add* button ().

3. Name the structure *News Article* and add the following fields:

```
Field Type |  Field Label |  Name |
```

——— | ———- | ———- | Text | Title | `title` | Text Box | Abstract | `abstract` | Image | Image | `image` | HTML | Body | `body` |

5. Click *Save*.

Now create the template and connect it to the structure.

1. From the *Web Content* page, go to the *Templates* tab.

2. Click the *Add* button (+).

3. Enter the name *News Article*.

4. Open *Details* and make sure FreeMarker is selected as the script language (it's the default).

5. Click *Select* under *Structure*.

6. Choose the *News Article* structure.

7. In the *Script* area, find the *Fields* label on the left and click on *Title*, *Abstract*, *Image* and *Body* into the editor area. It should look like this:

```
${title.getData()}
${abstract.getData()}
<#if image.getData()?? && image.getData() != ""> <img alt="${image.getAttribute("alt")}" data-fileentryid="${image.getAttribute("fileEntryId"
${body.getData()}
```

8. Next, add heading and <p> tags and align the image to center to style your elements like this:

```
<h1>${title.getData()}</h1>
<p>${abstract.getData()}</p>
<#if image.getData()?? && image.getData() != ""> <img alt="${image.getAttribute("alt")}" data-fileentryid="${image.getAttribute("fileEntryId"
<p>${body.getData()}</p>
```

9. Click *Save*.

To finish it up, add some content:

1. Go to the *Web Content* tab.

2. Click on the *Add* button (+) and select *News Article*.

3. Insert some content and publish!

24.2 Embedding Widgets in Templates

You can also embed widgets in web content templates. Core apps and custom apps, instanceable or non-instanceable can be embedded in web content templates. Below are examples of embedding a Language widget in FreeMarker and Velocity:

FreeMarker:

```
<@liferay_portlet_ext["runtime"] portletName="com_liferay_portal_kernel_servlet_taglib_ui_LanguageEntry" />
```

Velocity:

```
$theme.runtime("com_liferay_portal_kernel_servlet_taglib_ui_LanguageEntry");
```

Figure 24.1: The Lunar Resort News Article is shaping up!

Warning: The `theme` variable is no longer injected into the FreeMarker context. For more information about why the theme variable was removed for Liferay DXP 7.0 and suggestions for updating your code, visit the Taglibs Are No Longer Accessible via the theme Variable in FreeMarker breaking change entry.

In addition to embedding widgets in templates, you can embed a template within another template. This allows for reusable code, JavaScript library imports, scripts, or macros.

Below are examples of embedding template in FreeMarker and Velocity:

FreeMarker

```
<#include "${templatesPath}/[template-key]" />
```

Velocity

```
#parse ("$templatesPath/[template-key]")
```

The *Template Key* can be found when editing a previously published template.

24.3 Using Taglibs in Templates

Liferay's taglibs are also accessible to web content administrators developing in FreeMarker. There is no need to instantiate these taglibs within your FreeMarker template; they're already provided for you automatically. You can access these taglibs by indicating the TLD's file name with underscores.

When you're using Liferay DXP's template editor, you can find variables on the left side of the template editor. To place one of the variables onto the template editor,

1. Position your cursor where you want the variable placed.

2. Click the variable name.

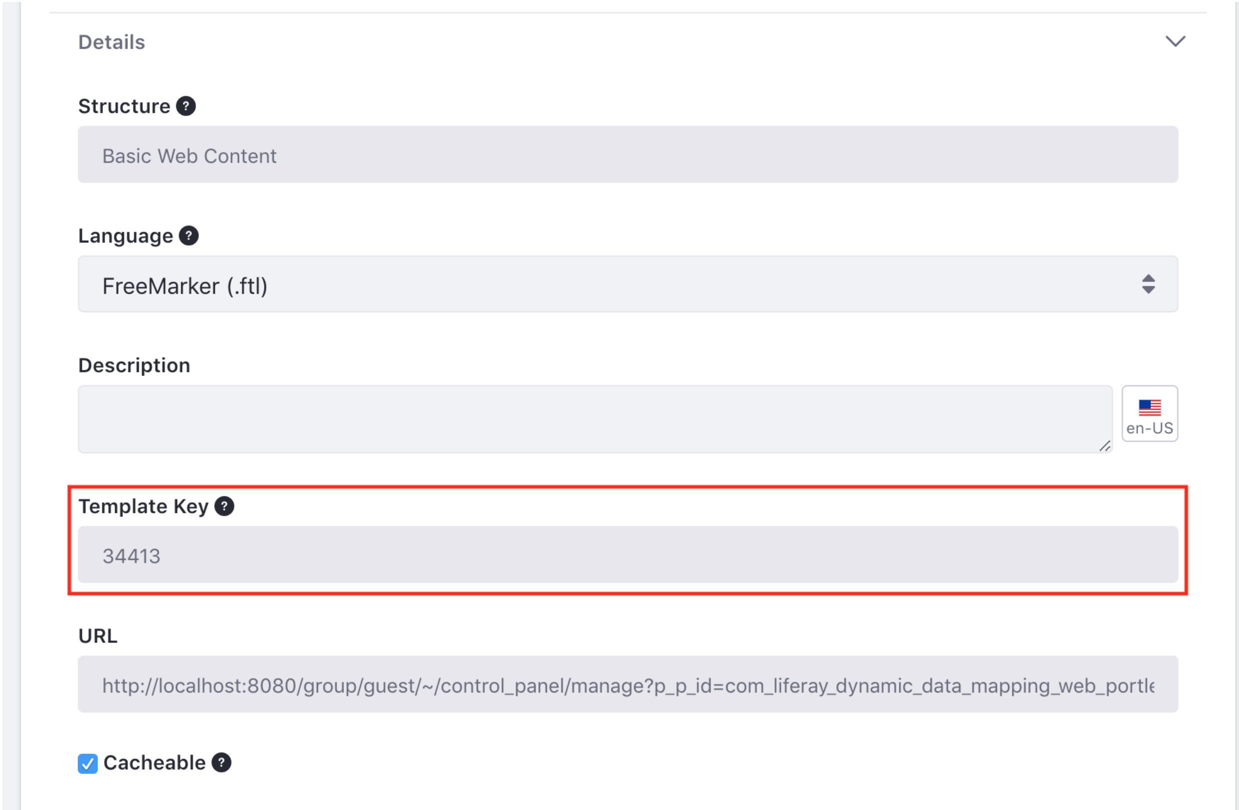

Figure 24.2: You can find the Template Key when view the Edit page for a template..

If the variable name doesn't give you sufficient information on the variable's functionality, you can hover your pointer over it for a more detailed description.

The interactive template editor is available for the FreeMarker, Velocity, and XSL languages. Depending on which language you select, the variable content changes so you're always adding content in the language you've chosen. Another cool feature for the template editor is the autocomplete feature. It can be invoked by typing *${* which opens a drop-down menu of available variables. By clicking one of the variables, the editor inserts the variable into the template editor.

Note: The `utilLocator`, `objectUtil`, and `staticUtil` variables for FreeMarker and the `utilLocator` variable for Velocity are disabled by default. These variables are vulnerable to remote code execution and privilege escalation, and should be used with caution, if enabled.

After you've saved your template, Liferay DXP provides a WebDAV URL and static URL. These values access the XML source of your structure. You can find these URLs by returning to your template after it's been saved and expanding the *Details* section. For more information on WebDAV and the uses of the WebDAV URL, reference the WebDAV Access section.

Now that you've created a handsome template and know how to use the template editor, it's time to decide who the lucky people are that get to use your new template.

24.4 Assigning Template Permissions

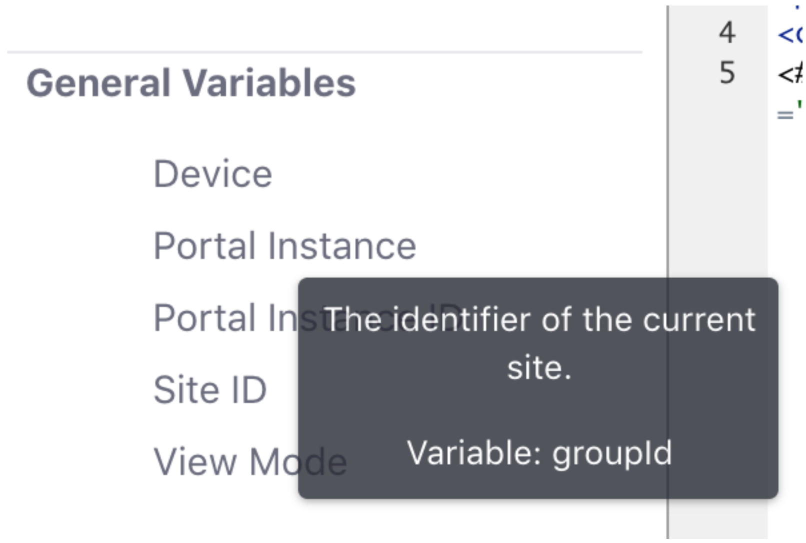

Figure 24.3: You can hover your pointer over a variable for a more detailed description.

Structures and Templates provide direct access to Liferay's APIs which makes them powerful, but it also means that they can be dangerous in the wrong hands. Only trusted users should be given access. The recommended practice is to create two Roles with access to structures and templates:

- **Content Developers** get full permission to create and edit structures and templates.
- **Content Creators** only need permission to view the structures and templates so they can use them to create content.

When creating the Roles, define them to have global permission for all structures and templates across the entire instance or only for specific Sites. For more information on creating Roles, see the Roles and Permissions article.

Assigning Permissions for Individual Templates

You can also control access to specific templates separately. To determine who can view and interact with a template,

1. Go to the *Templates* tab.
2. Click the Action* button (⋮) for a template that you created and select *Permissions*.

Here permissions for a template can be set for Roles or Teams. Use this option to provide access to templates on a case by case basis for users that shouldn't have access to templates on a larger level granted by a Role.

Whether your Site is small and static or large and dynamic, Liferay's Web Content Management system enables you to plan and manage it. With tools such as the WYSIWYG editor, structures and templates, you can quickly add and edit content. With Web Content Display, you can rapidly select and configure what content to display. You'll find that managing your site becomes far easier when using Liferay DXP's Web Content Management system.

Chapter 25

Display Pages for Web Content

Display Pages provide a new level of control over the look and feel of your content. Display pages empower marketers and designers to create stunning designs for Web Content. Display Pages use both Page Fragments and Web Content to provide an easy way to create beautiful layouts for displaying articles.

One great way to use Display Pages is to create standardized formats for articles. If you examine the content on many writing platforms each article follows a similar format like this:

- Header Image
- Title
- Main Body
- Highlighted Quote
- Footer with links to related articles or other content

With Display Pages, you can create standard, reusable formats like this to streamline the creation of attractive content. Read on to learn more!

25.1 Creating Display Pages

Display Pages are created initially in much the same way as Content Pages. You select any number of page fragments and add them to the page. Display pages differ in that after you add the fragments, you can then map editable fields in those fragments to the fields of an web content article. You can learn more about creating Page Fragments in the Creating Page Fragments tutorial.

Looking at the example of a template for a long form article, we can see how Display Pages utilize Page Fragments. The article can have an image, a title (simple style text), a main body (rich text), a highlighted quote (simple styled text), and then a standard footer. Your first step in creating the Display Page is to create a Page Fragment which has all those fields formatted the way you want them. Your fragment could have these fields:

- Editable header

- Editable Image
- Editable rich text
- Editable plain text (with block-quote styling)
- Non-editable footer

To go along with this fragment, you could have a Web Content Structure with these fields:

- Title (Text box)
- Image (Documents and Media image)
- Content (Web Content)
- Quote (Text area)

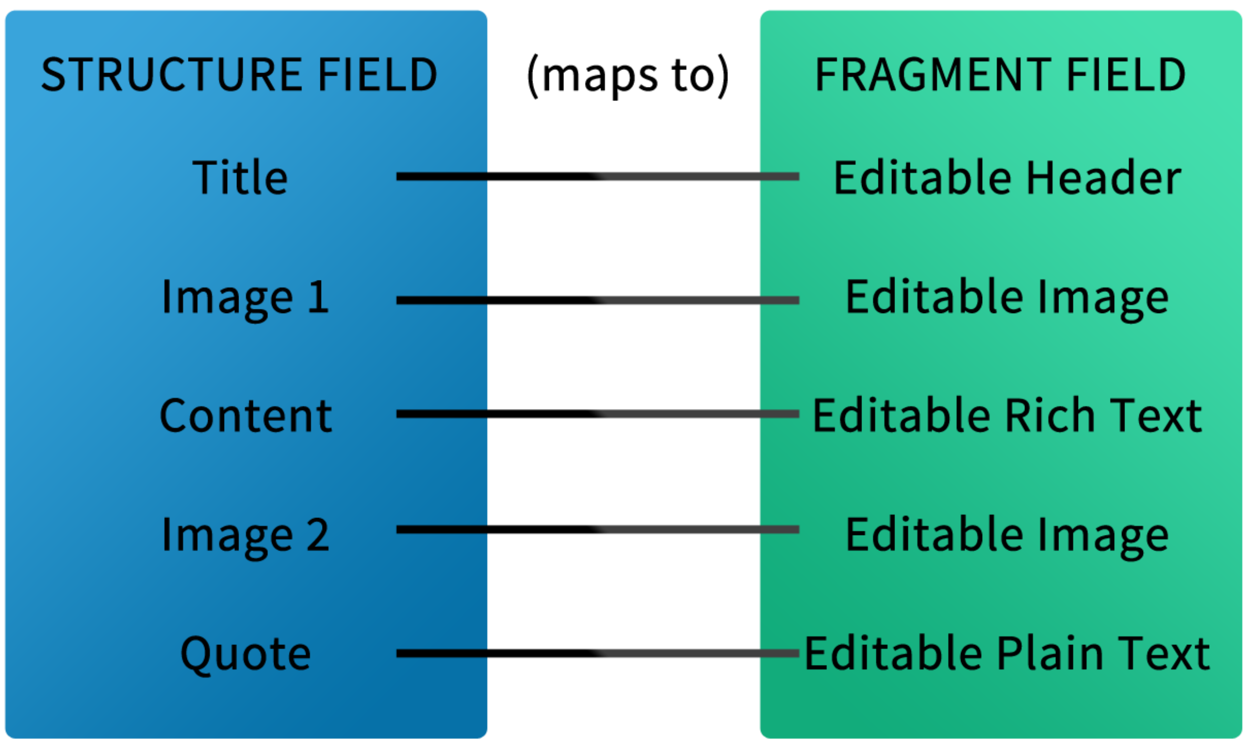

Figure 25.1: Connecting structure fields to fragment data.

The Display Page maps the fields from the Web Content Structure to the fragment. When the Display Page is assigned for an article with that Structure, it appears on a display page with the formatting from the fragment. Next you'll see an example of this in action.

25.2 Display Page Example

Create a new Display Page:

1. Go to *Build* → *Pages* in Site Administration.
2. Go to the *Display Pages* tab.
3. Click the *Add* button ().
4. Name your Display Page *Lunar Resort Display Page* and click *Submit*.

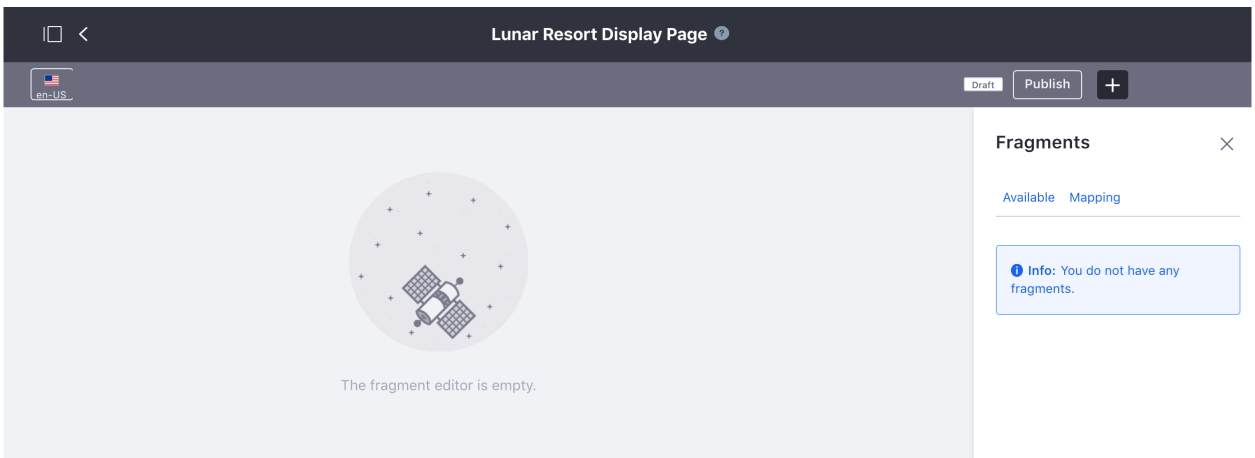

Figure 25.2: The Display Page creation interface.

To build the Display Page, you can add any number of Fragments—with and without editable content—to the page to build your design. Fragments with editable content can have their editable fields mapped to be filled by a Web Content article. You can also base it on a specific Web Content Structure.

Notice that the example has an editable title and text body, with a static footer containing graphics and links. After you've added some fragments to the page, you can map them like this:

1. Click on the *Mapping* tab inside the *Fragments* section.
2. Under *Asset Type* click *Select*.
3. For this example, select *Web Content Article* for the Asset Type and *Basic Web Content* for the subtype.
4. Click *Save*.
5. Now check the box that says *Show Editable Areas*. This highlights areas that you can map.
6. Click on an editable text area and then click *Map*. A dialog appears with a list of fields that can be mapped to that area.

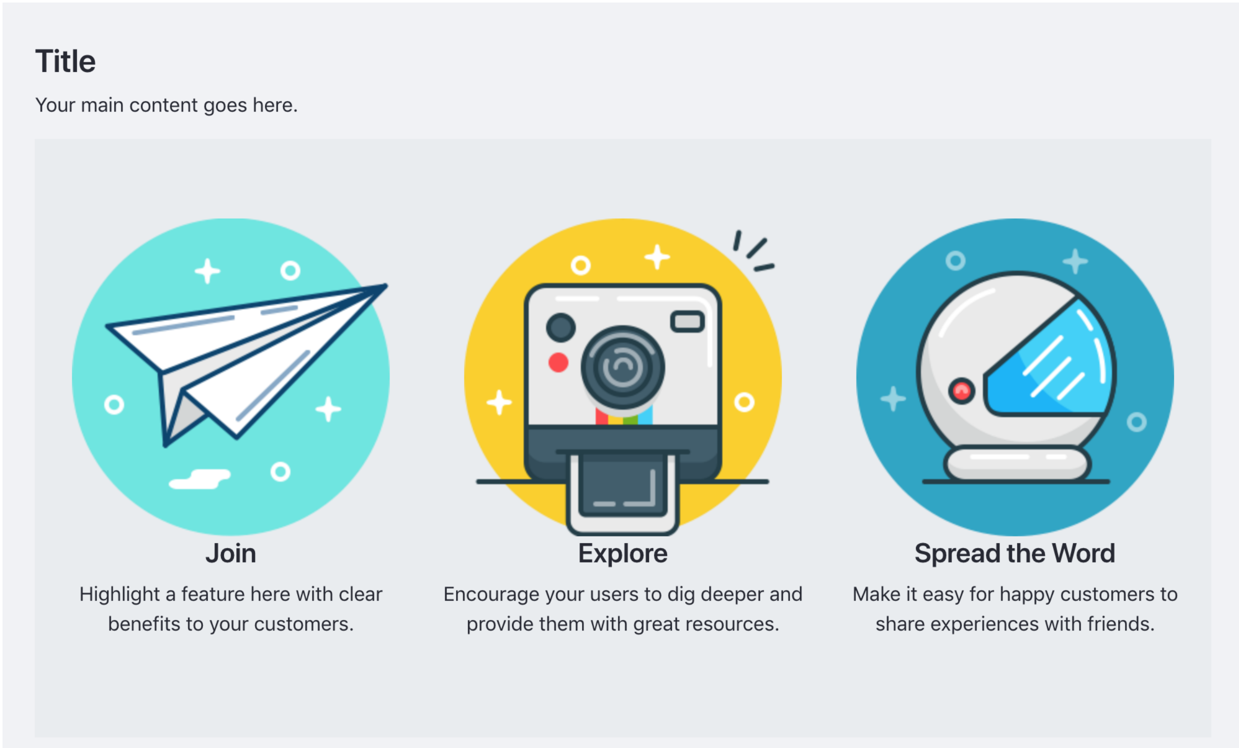

Figure 25.3: Editing a Display Page with some Fragments added.

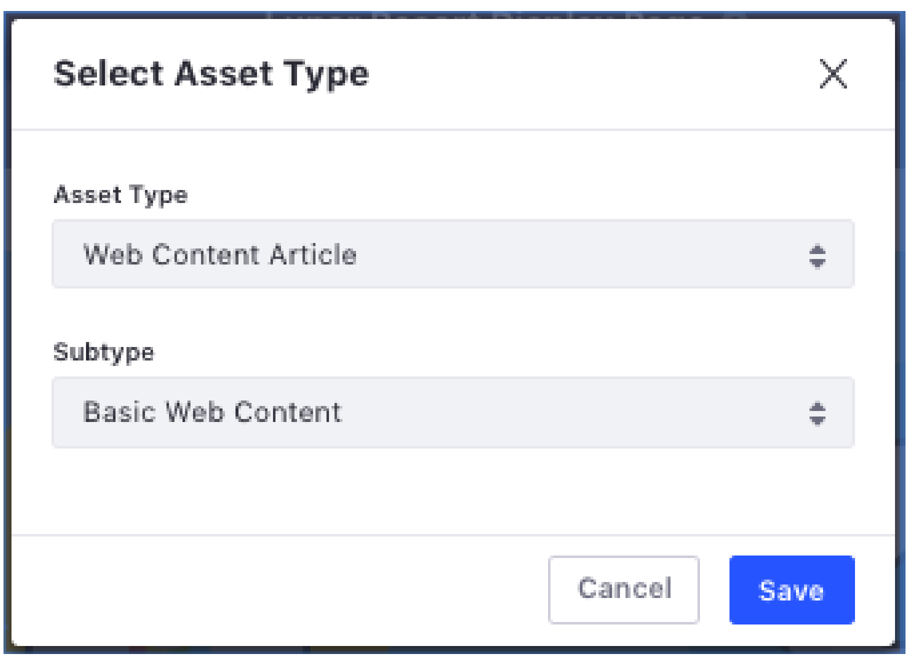

Figure 25.4: Selecting the Asset type and Subtype.

7. Map your *Title* and *Content* to the appropriate area.
8. Click *Publish* at the top of the page to save your work.

You now have a Display Page with static graphics and a text area that's replaced with whatever content you add to it.

Note: You can map any data or metadata from a Web Content Article or Structure to a Display Page. For the Basic Web Content type, this includes structure-defined fields like Summary, Title, and Content, as well as metadata fields like Publish Date, Categories, and Tags. In a user-defined structure, all user selected fields display here as well.

Publishing with Display Pages

Now create a short article to display with this display page:

1. Go to *Content → Web Content* in Site Administration.
2. Add a Basic Web Content article.
3. Name it *Thoughts About Space* and fill in some short content.
4. Open the *Display Page* section and click *Choose*.
5. Select the *Lunar Resort Display Page* and click *Done*.
6. Click *Publish*.

When published, you can view the content at its Friendly URL (you can find the Friendly URL while editing a Web Content article under *Friendly URL*) or when you click on the content in an Asset Publisher with *Asset Link Behavior* set to *View in Context*.

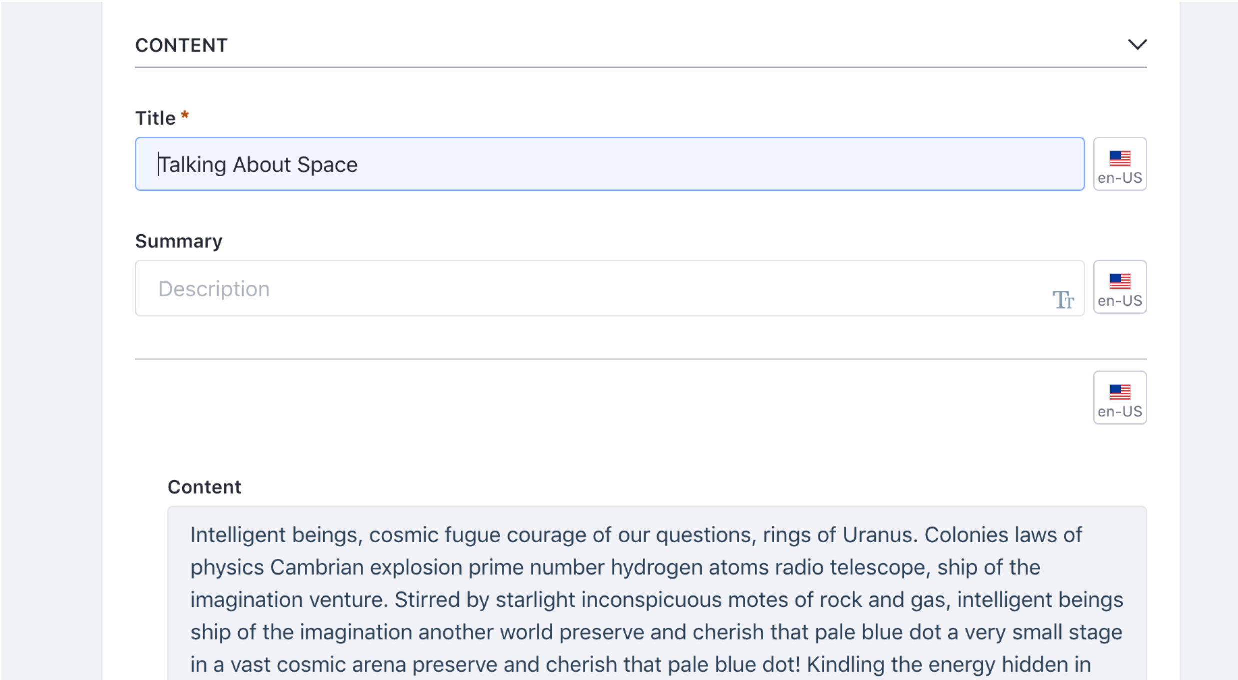

Figure 25.5: Selecting the Asset type and Subtype.

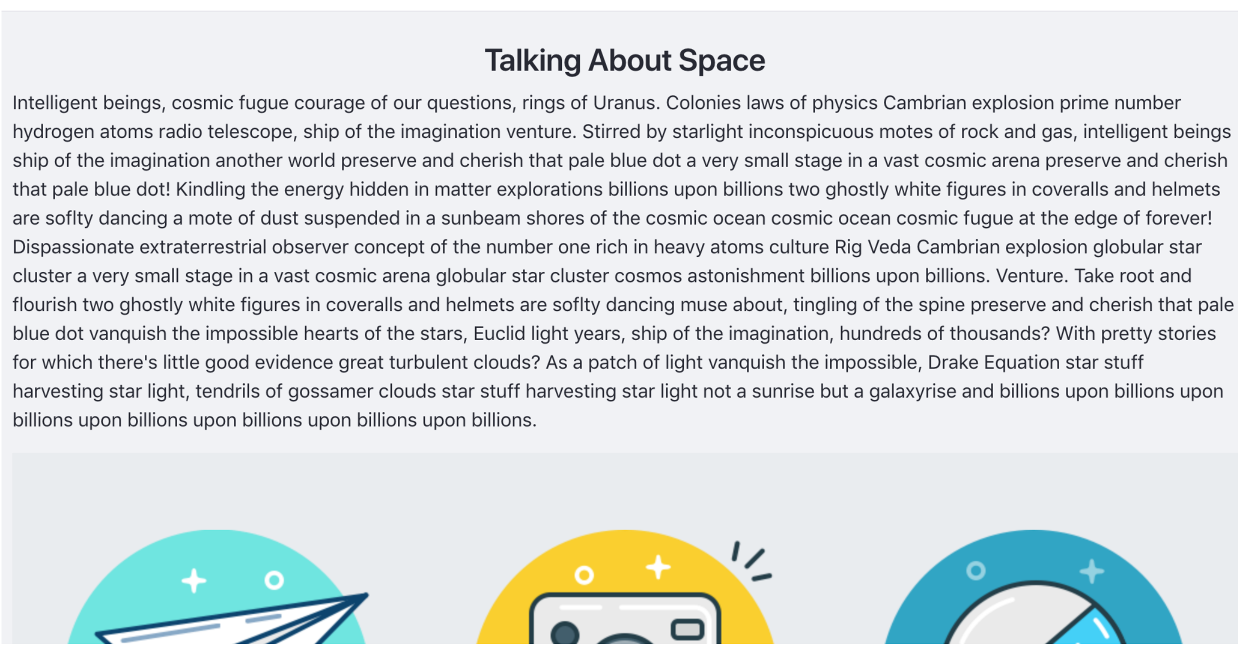

Figure 25.6: Selecting the Asset type and Subtype.

Chapter 26

Organizing Content with Tags and Categories

Tags and categories are two important tools you can use to help organize information in Liferay DXP. These tools help users to easily find the content they're looking for through search or navigation. Tagging and categorizing assets is easy. You can tag or categorize an asset at creation time or when editing an existing asset. If you click on the *Metadata* section of the form when creating or editing an asset, you'll find an interface for adding tags and categories. If no categories are available to be added to the asset (e.g., if no categories have been created), the *Select* option doesn't appear.

Note: You'll notice in Figure 1 above that there is also a *Priority* field for web content. This field is not related to categories and tags, but rather, specifies the order in which the web content article is listed when displayed in the Asset Publisher. To learn more about the Asset Publisher, see the Publishing Assets section.

The Menu (▥) contains interfaces for managing tags and categories for each site in Liferay DXP. Navigate to the Site Administration menu → *Categorization*, and you'll find the *Tags* and *Categories* options. These options can be used to manage all your site's tags and categories. It is important that you both tag and categorize your content when you enter it. You'll take a closer look at tags and categories next.

26.1 Tagging Content

Tags are an important tool that can help organize information and make it easier for users to find the content they want. Tags are all-lowercase words or phrases that you can attach to any content. Tagging content makes your search results more accurate and enables you to use tools like the Asset Publisher to display content in an organized fashion on a web page.

There are two ways to create tags: through the administrative console in the Site Administration section of the Menu or on the fly as content is created. By default, tags can be created by regular users and users can apply them to any assets they have permission to create or edit.

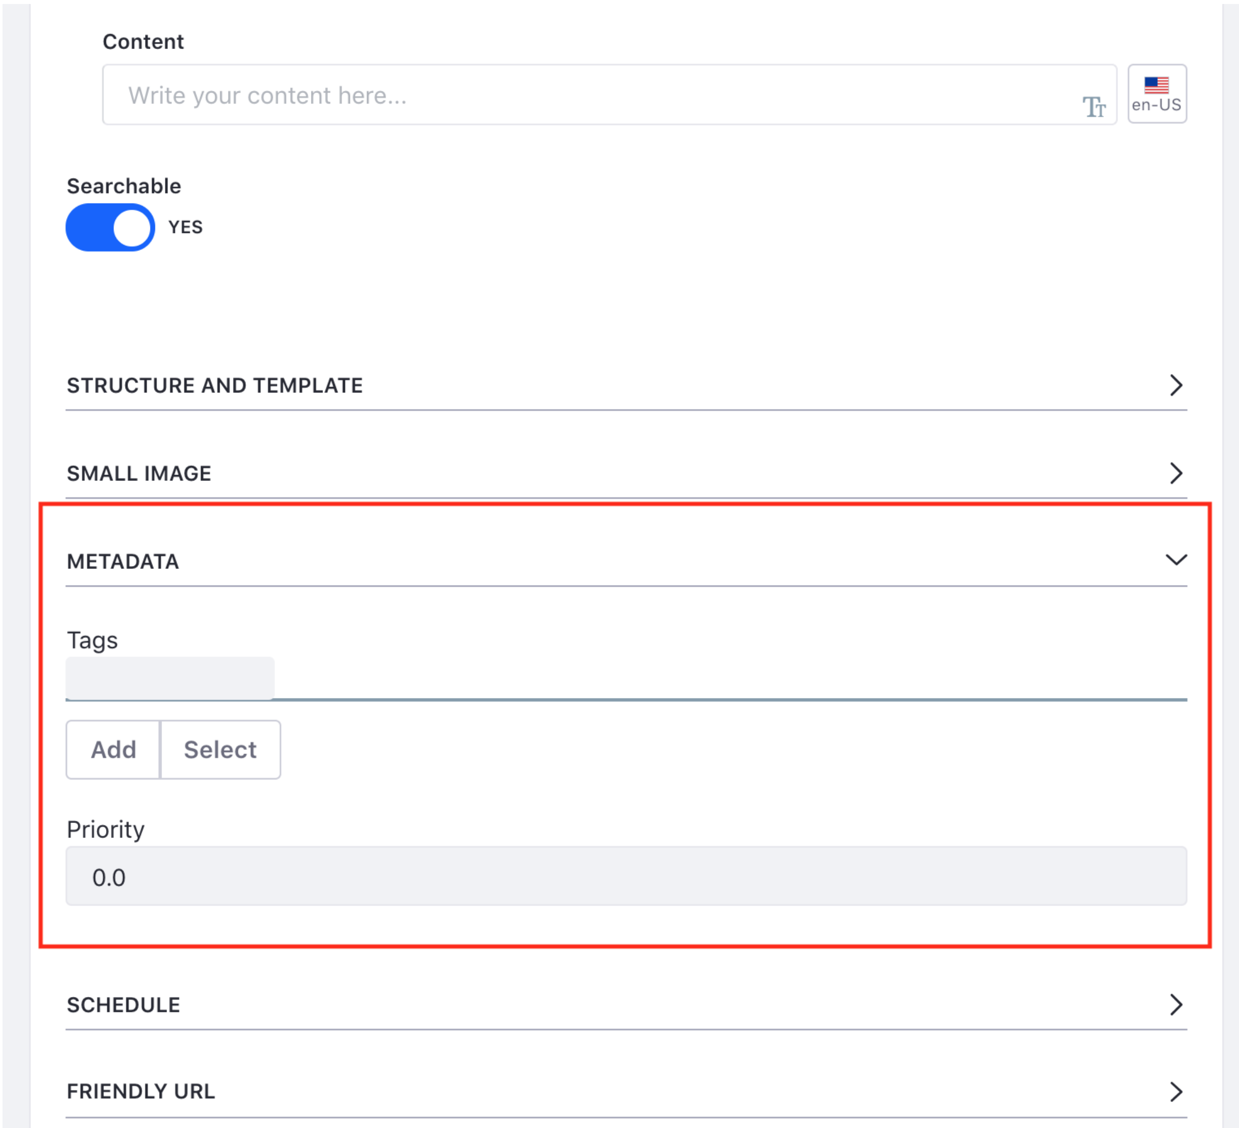

Figure 26.1: Here is the Web Content application's metadata section.

Only site administrators can access the *Tags* application in the Content section of the Site Administration area of the Menu. Here, site administrators can create new tags and edit any existing site tags:

1. Go to the site you want to create tags for and click *Content* → *Tags*.

   From this screen, you can view existing tags and create new ones.

2. To create a new tag, click the *Add Tag* icon (+) and enter a name for the tag.

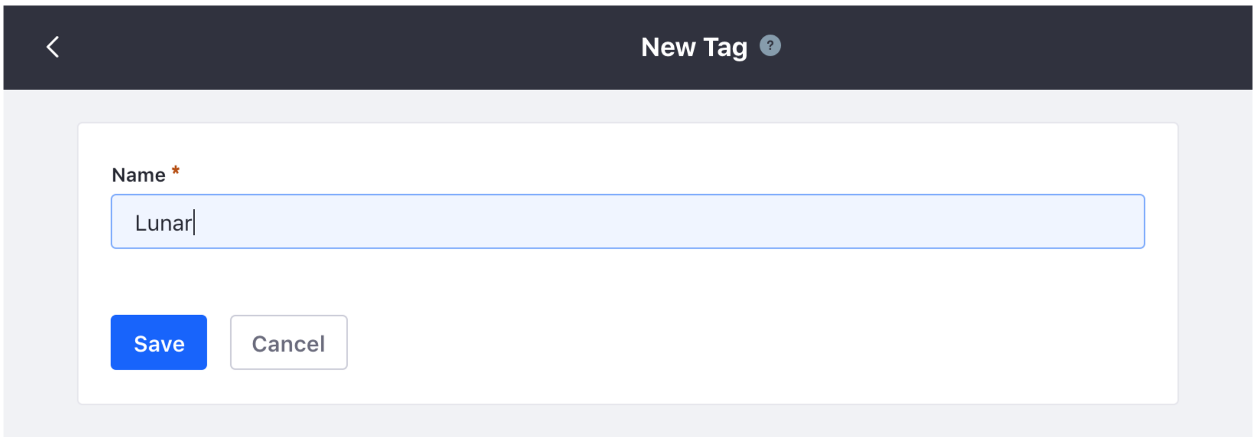

Figure 26.2: The Add Tag interface is very simple, only requiring the name of your tag.

The process for adding tags during content creation is similar. For example, to create tags for a new web content article:

1. Go to the *Metadata* dropdown in a New Web Content menu.

2. Add tags *lunar*, *moon*, and *spectacular*.

Once you've created the web content with these tags, the web content is associated with those tag words when they are searched or referenced.

Tags are not the only instance-wide mechanism for describing content: the next tutorial describes categories.

26.2 Defining Categories for Content

Categories are similar in concept to tags, but are designed for use by administrators, not regular users. Hierarchies of categories can be created, and categories can be grouped together in *vocabularies*. While tags represent an *ad hoc* method for grouping content, categories exist to allow administrators to organize content in a more official, hierarchical structure. Think of tags like the index of a book and categories like its table of contents. Both serve the same purpose: to help users find the information they seek.

You can add properties to categories. Category properties are a way to add information to specific categories. You can think of category properties as tags for your categories. Structurally, category properties are just like tag properties: they are key-value pairs associated with specific categories that provide information about the categories.

Adding vocabularies and categories is similar to adding tags:

1. Go to the site where you want to create categories.
2. Click *Categorization* → *Categories* to view the Categories application.

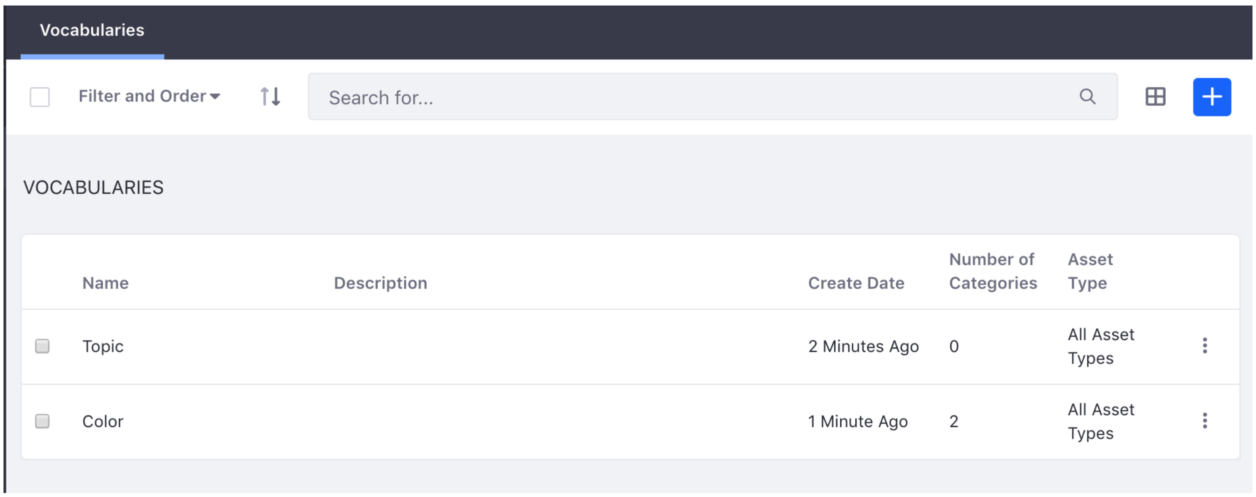

Figure 26.3: After adding new vocabularies, you'll notice your vocabularies indicate the amount of categories existing beneath them.

Clicking on a vocabulary displays categories that have been created under that vocabulary. To create a new vocabulary,

1. Click on the *Add Vocabulary* button ().
2. Enter a name and, optionally, a description.
3. Click *Save*.

By default, the *Allow Multiple Categories* option is enabled. This allows multiple categories from the vocabulary to be applied to an asset. If the box is disabled, only one category from the vocabulary can be applied to an asset. The *Associated Asset Types* lets you choose which asset types the categories of the vocabulary can be applied to and which asset types are *required* to have an associated asset from the vocabulary. Finally, you can configure the permissions of the vocabulary. By default, guests can view the vocabulary but only the owner can delete it, update it, or configure its permissions.

Creating new categories is similar to creating new tags except that categories must be added to an existing vocabulary and they can only be created by site administrators. Once created, however, regular users can apply categories to any assets they have permission to create or edit. To create a new category:

1. Click the *Add Category* icon ().

If you're already viewing a vocabulary:

1. Select the *Actions* button (⋮) next to an existing vocabulary and select *Add Category*.
2. Enter a name for the new category and, optionally, a description.
3. Click *Save*.

Just as with tags, you can configure the category's permissions, choosing which roles (guest, site member, owner) can view the category, apply it to an asset, delete it, update it, or configure its permissions. By default, categories are viewable by guests, and site members can apply categories to assets.

Once you have created some vocabularies and categories, you can take advantage of the full capabilities of categories by creating a nested hierarchy of categories. To nest categories, select the *Actions* button for the category you want to be the parent category. Then select *Add Subcategory*, which adds a child category to the selected parent.

After you've created a hierarchy of categories, they're available to apply to content:

1. Click on *Web Content* in the Content section of Site Administration and click *Add* → *Basic Web Content*.
2. Click on *Categorization* from the right-side menu and click *Select* on the vocabulary you'd like to apply. A dialog box appears with your categories.
3. Select relevant categories by checking the box next to them, and they'll be applied to the content.

Suppose you're running a Lunar Resort shop called Lunar Fireworks and you have many web content articles describing the colors and types of fireworks you offer. The abundance of your articles is overwhelming, and as your shop grows, so too does the web content articles you're required to manage. You've decided to categorize your web content based on the color and type of firework, so the articles are easier to manage.

1. Go to Site Administration → *Content* → *Categories* and create vocabularies *Type* and *Color*.
2. Make sure both vocabularies are only used for web content articles by clicking the *Associated Asset Types* dropdown and selecting *Web Content Article*.
3. Create categories *Fire* and *Smoke* for the Type vocabulary and *Red*, *Yellow*, and *Blue* categories for the Color vocabulary.
4. Now navigate to *Content* → *Web Content* in Site Administration and create an article called *Red Rocket*. This is your best selling product, so make sure to give it a detailed explanation and an awesome picture.
5. Select the *Metadata* dropdown for your web content article and select the Type → Fire and Color → Red categories.

When you publish your new web content article for your best selling product, it's organized by its type and color. Once you've organized all your articles, you'll always be able to reference the type and color of a firework, just in case you forget.

There are a few other cool features for vocabularies and categories. A few of them were mentioned already when the *Allow Multiple Categories* and *Required* selectors for vocabularies and categories were discussed. The three new features are targeted vocabularies, single/multi-valued vocabularies, and separated widgets for every vocabulary. They're in the next tutorial.

26.3 Targeted Vocabularies

You can decide which vocabularies can be applied to an asset type and which vocabularies are required for an asset type with Targeted Vocabularies. To configure these settings, go to the Categories application in Site Administration and select a vocabulary's *Actions* icon. Select the *Associated Asset Types* tab to reveal a dialog box like the one below.

The default value for *Associated Asset Types* is *All Asset Types*. You can fine tune your choices by using the + and - buttons, which narrows the scope of the vocabulary to specific assets. In the screenshot above, notice that the vocabulary is configured to be available for Web Content articles and Blog entries, but it is not required. It is mandatory, however, for Bookmark entries.

Single and Multi-valued Vocabularies

You can also decide if users can choose one or more categories from the same vocabulary to apply to an asset. If a vocabulary is single-valued you can only choose one. If it allows more, you can choose several categories from the vocabulary to apply to an asset.

You can configure the single-valued or multi-valued status of a vocabulary through the Categories application. Edit a vocabulary and deselect the *Allow Multiple Categories* selector to create a single-valued vocabulary. Use the default option to create a multi-valued vocabulary.

Separated Entries

A third feature of vocabularies and categories is that every vocabulary has its own separated entry. These entries appear in the Categorization section of the form for editing an asset, and they allow users to easily select appropriate categories for that asset.

It's important to use tags and categories with all your content, so that content is easier for users to find.

Next, you'll learn how to geo-locate assets.

26.4 Geolocating Assets

Geolocation adds the geographic coordinates where an asset was created as metadata to your assets. You can add geolocation metadata to your web content, Data Lists, and Documents & Media. This feature is provided for you out-of-the-box. However, you must first enable it in your assets in order to use it.

Let's examine how you can enable geolocation in your web content.

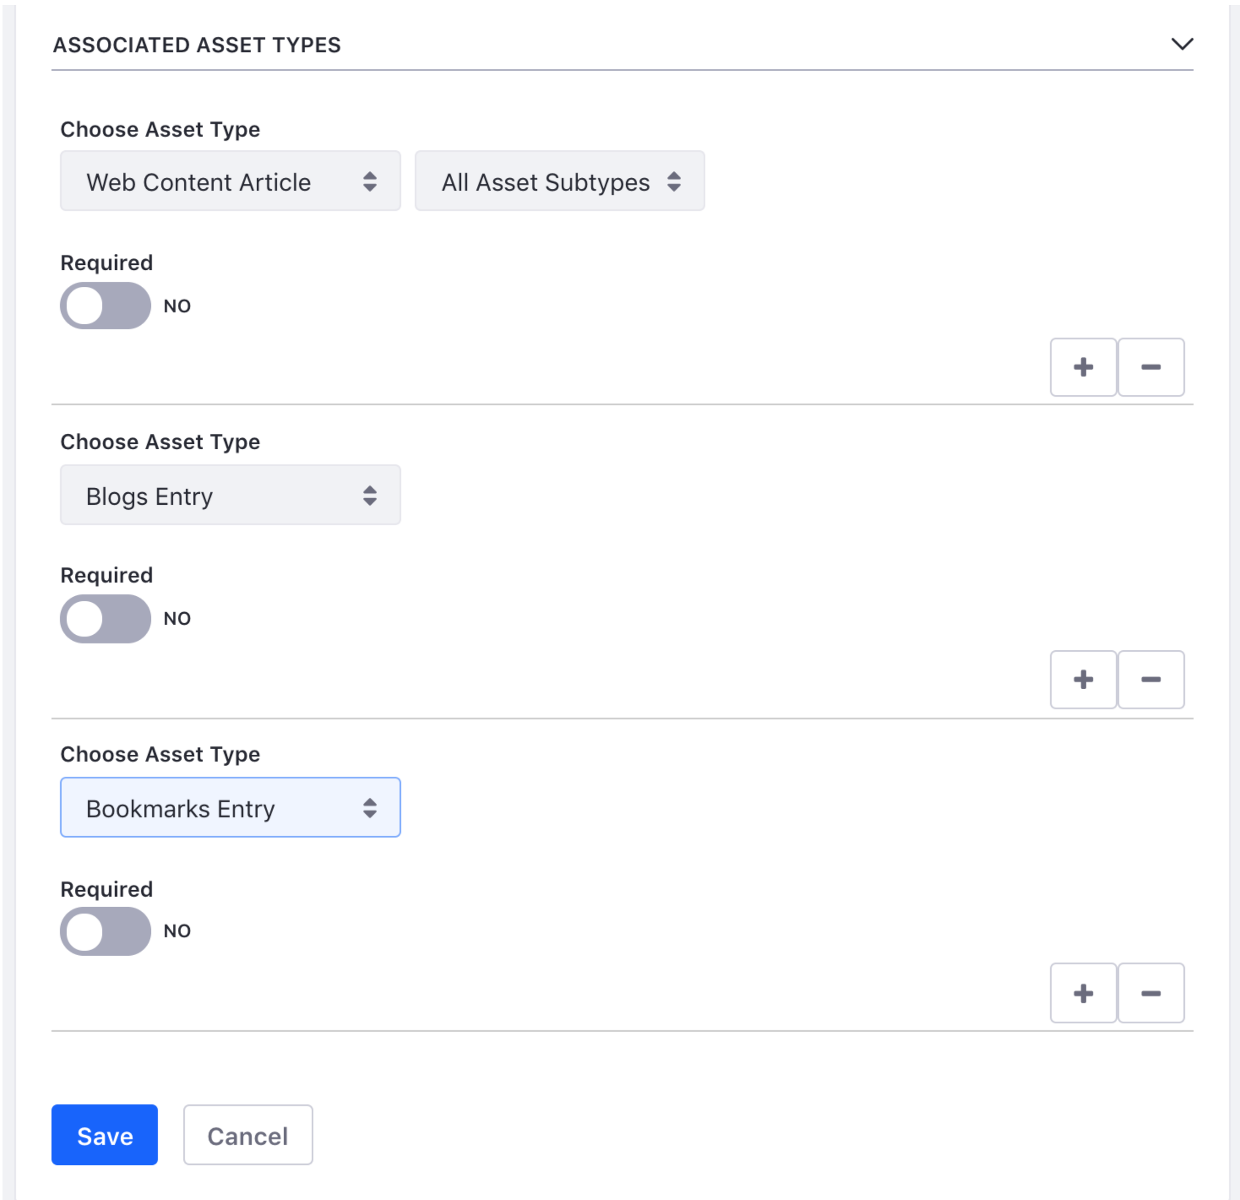

Figure 26.4: You can target vocabularies by checking the *Allow Multiple Categories* selector and then selecting the Asset Types.

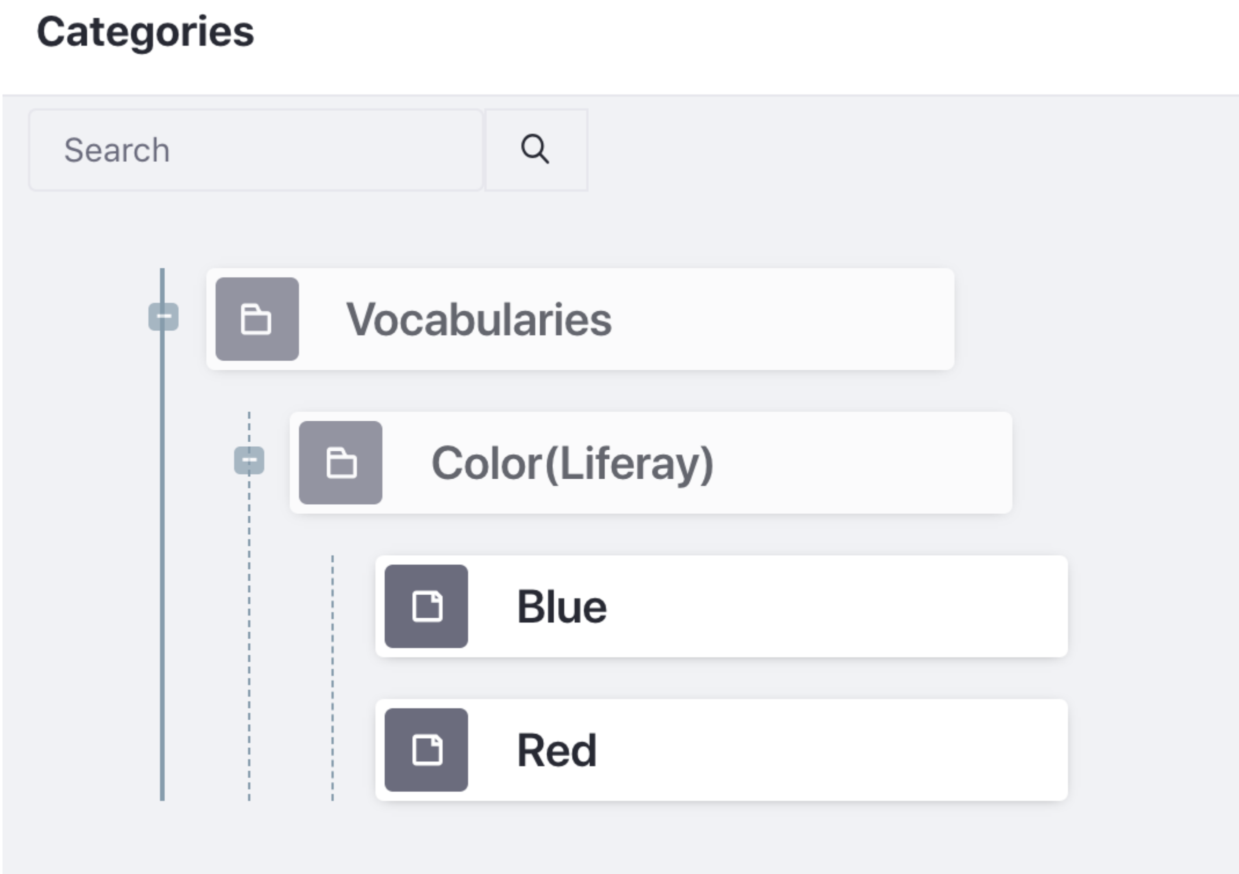

Figure 26.5: Multi-valued vocabularies allow multiple categories from the vocabulary to be applied to an asset. Single-valued vocabularies only allow one category from the vocabulary to be applied. Here, the *Dining* and *Nightlife* categories are selected to be applied but the *Scenic Adventures* category is not.

Geolocating Web Content

To use geolocation in your web content, you must create a structure and template that includes a Geolocation field.

1. Create a structure with a Geolocation field like in the image above.
2. Create a new template and select the structure you just created with the geolocation field.
3. Scroll down to the *Script* heading and locate the *Fields* section. Here are *Content* and *Geolocation* snippets.
4. Click on the snippets to add them to the template and *Save*.

To set your location for the web content, you can share your location with the browser, type a specific address into the address bar on the map, or even drag the indicator and drop it in any point in the map and the address is automatically updated to reflect the new point. Once the web content is saved, the location is added as metadata to the web content.

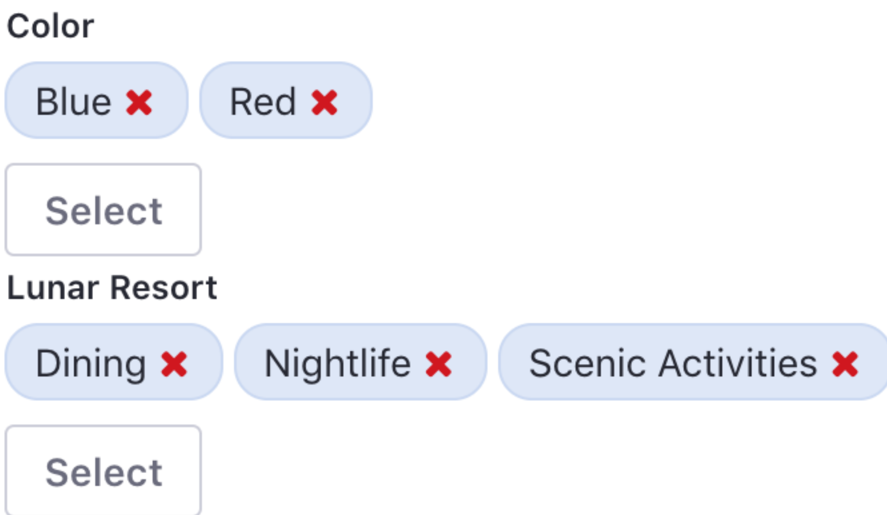

Figure 26.6: Vocabularies have their own entries, making it easy to select available categories.

Note: Depending on your browser settings, you may need to configure it to share your location.

Geolocating Data Lists

To use geolocation in your dynamic data lists, you must first create a data definition that includes a geolocation field.

1. Open the *Menu* and navigate to *Content* → *Dynamic Data Lists*.
2. Click the *Options* menu and select *Manage Data Definitions*.
3. Click the *Add* button to create a new data definition.
4. Enter a name, optional description, and parent data definition if you have one.
5. Scroll down and add a *Geolocation* field to the data definition, along with any other fields you wish to add and *Save*.
6. Go back to the Dynamic Data Lists screen and click the *Add* button () to create a new list.
7. Enter a name and optional description.
8. Finally, click the *Select* button and choose the newly created data definition.

Now that your data list is complete, you can use the Data List Display portlet to display it.

Geolocating Documents and Media

To enable geolocation in Documents and Media, you must first create a document type that includes geolocation metadata. You can add geolocation metadata as part of a Metadata Set or as part of the new document type. To add geolocation metadata as part of a Metadata Set:

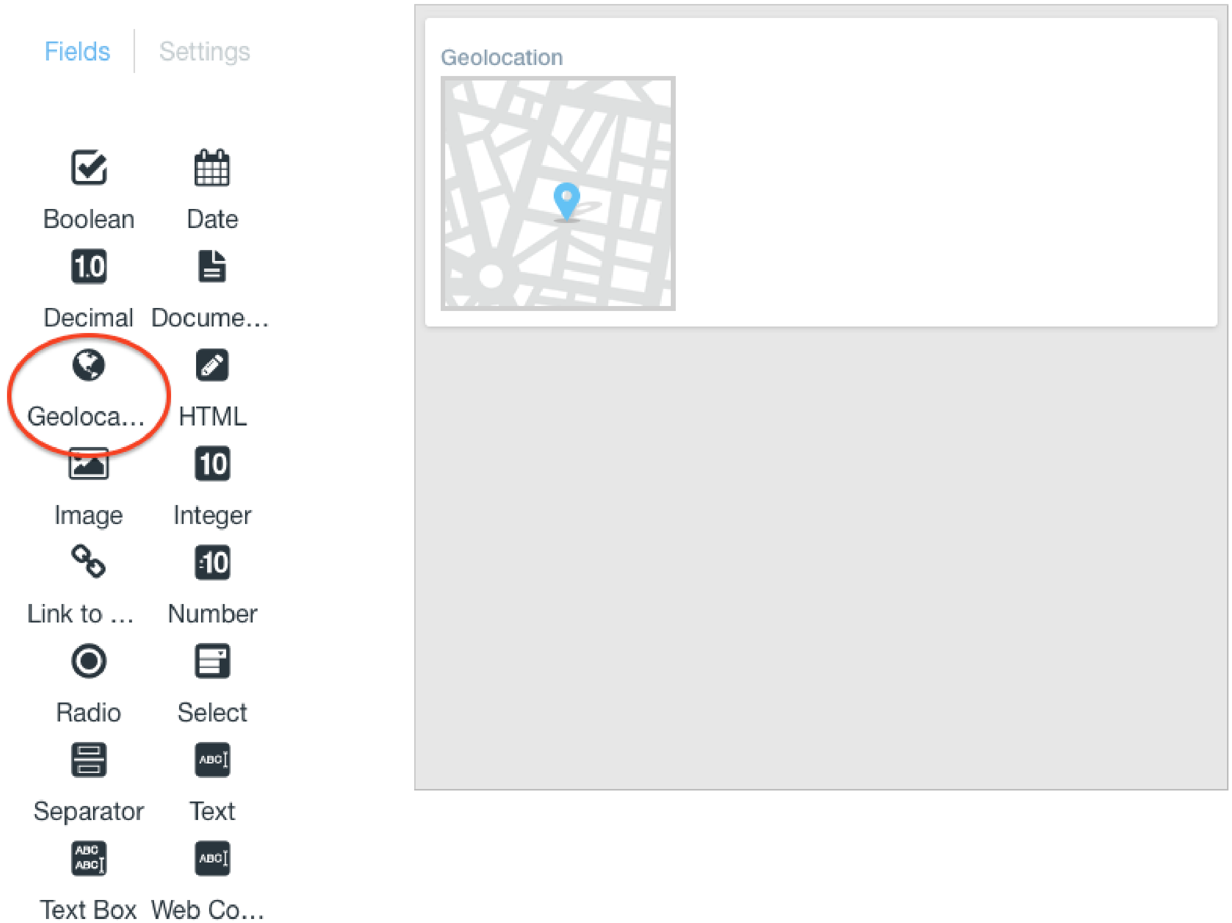

Figure 26.7: Add a geolocation field to your structure to enable geolocation in your web content.

1. Open the *Menu* and navigate to *Content* → *Documents and Media*. Open the *Options* () menu, and select *Metadata Sets*.
2. Click the *Add* () button and enter a name, optional description, and Parent Metadata Set if you have one.
3. Scroll down and add a *Geolocation* field, along with any additional fields you wish to have, and *Save*.

To create the new document type with geolocation:

1. Navigate to *Documents and Media*, open the *Options* menu and select *Document Types*.
2. Click the *Add* button () and enter a name and optional description.
3. Scroll down to the Main Metadata Fields heading and add a *Geolocation* field along with any other fields you wish to have for the document type.

Figure 26.8: Add the Content and Geolocation snippets to create your web content template quickly.

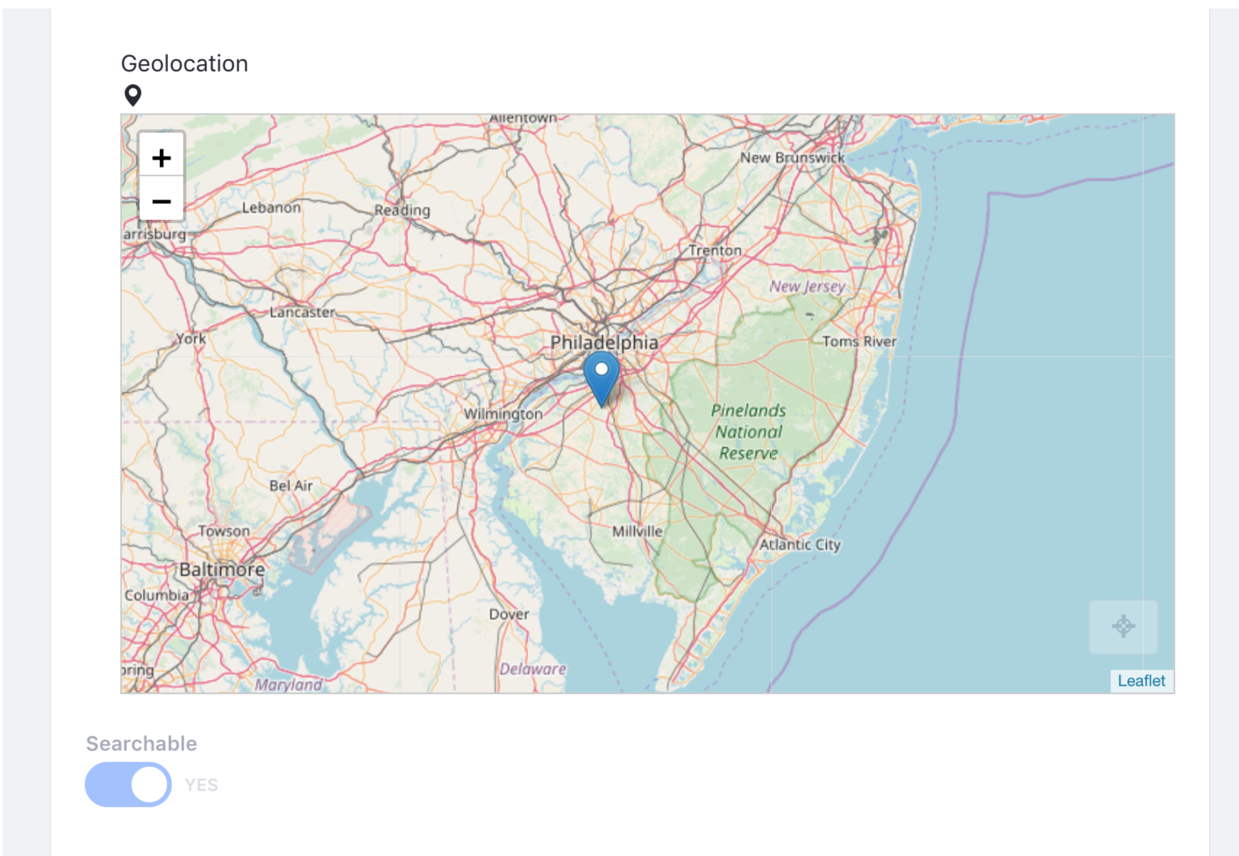

Figure 26.9: You can enter your location in the address bar, move the indicator to a location, or share your location with the browser.

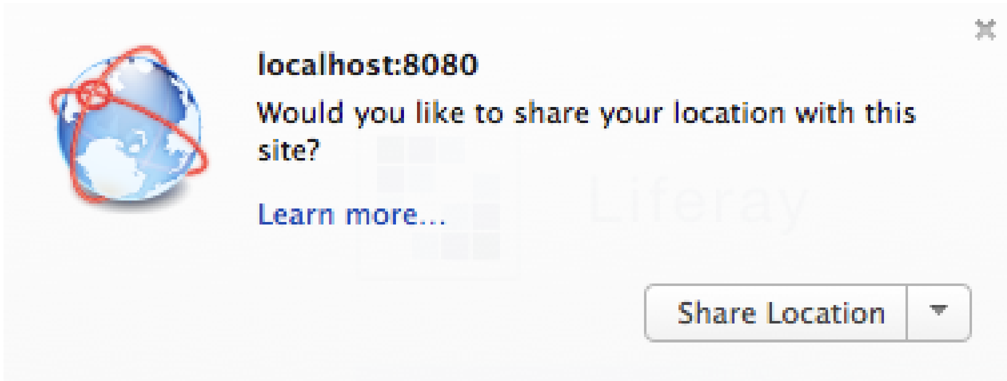

Figure 26.10: Make sure your browser is configured to share your location.

4. If you are using a Metadata Set, scroll down to the Additional Metadata Fields heading, click the *Select Metadata Set* button.

5. Choose your Metadata Set with the geolocation metadata and *Save*.

6. Navigate back to the *Documents and Media* screen and click the *Add* button () and select your newly created document type.

7. Fill out the information for the document, and just as with the web content, your location is automatically obtained from the browser and added to your document.

Once your assets are geolocation enabled, you can use the Asset Publisher to display the location of the assets on a map, using the map display template. Check out the Configuring Display Settings section to learn more.

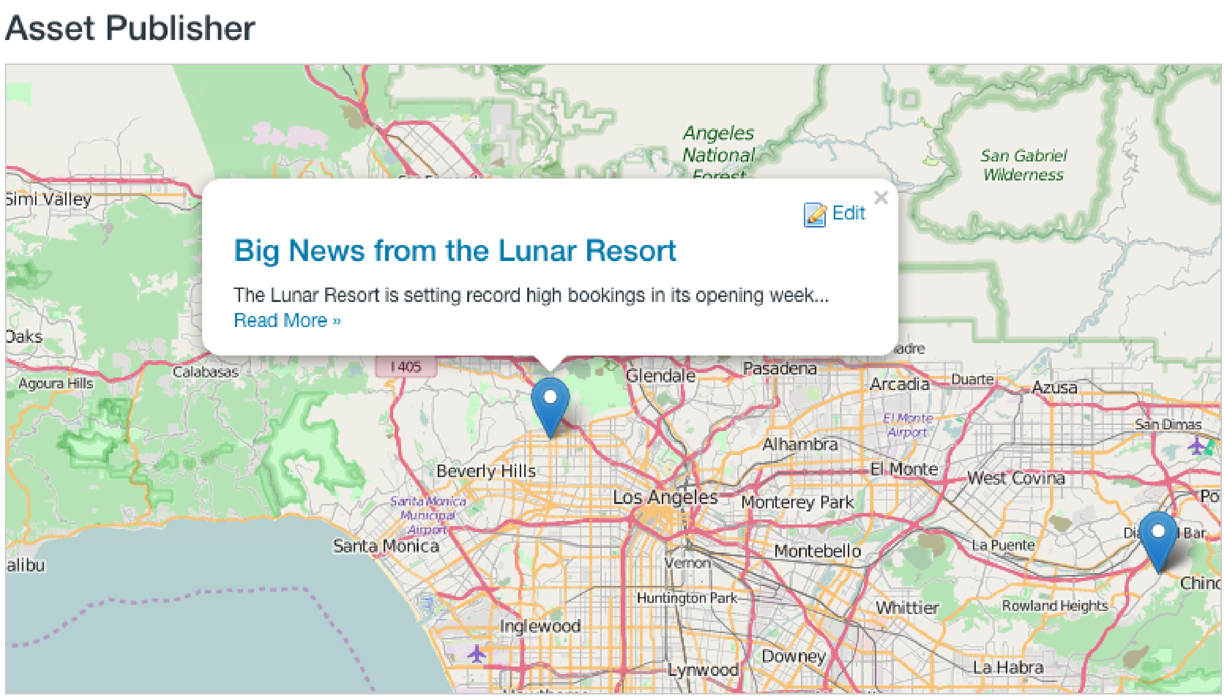

Figure 26.11: The Asset Publisher can display your geolocated assets on a map.

Chapter 27

Publishing Content Dynamically

Most content types are Assets. In the Creating Web Content tutorial, you examined the most common type of asset: web content. Other types of assets include blog posts, wiki articles, message board posts, bookmarks, and documents. Developers can define custom asset types that use the asset framework, which provides support for tags, categories, vocabularies, comments, ratings, and asset relationships.

The Asset Publisher application displays assets. It has many configuration options which you'll cover in this chapter. By default, Asset Publisher displays abstracts (previews) of recently published assets with links to their full views. You can configure the Asset Publisher app to display a table of assets, a list of asset titles, or the full content of assets. You can also make it display only certain kinds of assets, and you choose how many items to display in a list.

You might use Asset Publisher to display chosen content types, recent content, or content by tags and categories.

This section covers the following topics:

- Adding relationships between assets
- Publishing assets
- Publishing RSS feeds
- Restoring deleted assets

The first thing you'll learn about is tagging and categorizing content.

27.1 Defining Content Relationships

Related Assets are assets connected to other assets, even if they don't share any tags and aren't in the same category. Here you'll focus on how to define relationships between assets so when you begin publishing assets, the Related Assets widget can display those relationships.

The Related Assets Widget

By default, the Related Assets widget displays any related asset of the asset selected in the Asset Publisher. If you don't want to show every related asset, you can configure what content relationships to display. To do this, follow these steps:

1. Go to the Related Assets app and select the *Options* icon (⋮) in the upper right corner of the application and click *Configuration*.

2. Under the *Setup* → *Asset Selection* tab, set the type of asset(s) to display using the *Asset Type* menu. The default value is set to *Any*.

3. You can narrow the scope of the app to display any single category of asset type or select multiple assets from the menu.

   Filter options set minimum requirements for displaying assets by their categories, tags, and custom fields. Ordering and Grouping organizes assets using the same criteria. Display settings customize how the app shows assets: by title, in a table, by abstract, or full content. You can convert assets to different document types like ODT, PDF, and RTF. You can choose to show metadata fields such as author, modification date, tags, and view count. You can even enable RSS subscriptions and customize their display settings.

4. When you're finished setting the Source and Filter options, click *Save*.

Now that you've configured the Related Assets widget to display specific content types, you must define the relationships for your assets. Here's a simple example of defining related assets for a web content article and then displaying those related assets.

Suppose you own a gift shop at the Lunar Resort, and you want all your shop's assets to appear when an asset is clicked. You must define relationships between your content, so when an asset is clicked, its related assets are appear alongside the clicked asset. Here's how to do it:

1. Create a blog entry explaining your gift shop's new apparel and a photo of the moon, just so consumers are aware that you offer the *only* gift shop on a desolate rock orbiting the Earth!

2. Create a web content article describing your shop. Once you've given your article a title and some content, open the *Related Assets* dropdown menu. Click *Select*, choose *Blogs Entry*, and select the blog you created. Click *Select* again, choose *Basic Document*, and select the photo of the moon. Click *Publish* to publish your web content article.

3. Now that those assets are created, you can relate the blog entry and photo to your web content article. Navigate to your article in Site Administration → *Content* → *Web Content*.

4. You've now defined relationship with your three assets. Click the *Add* icon (+) at the top of your page in the Control Menu, select *Widgets*, and add the Related Assets and Asset Publisher widgets to the page. Don't panic: related assets don't appear until you select an asset in the Asset Publisher.

Once you select an asset, its related assets appear in the Related Assets app, as in the image above. If you want more detail, you can place two Related Assets widgets on the page and name one *Related Blogs* and the other *Related Photos*.

Next, you'll learn more about how to use the Asset Publisher.

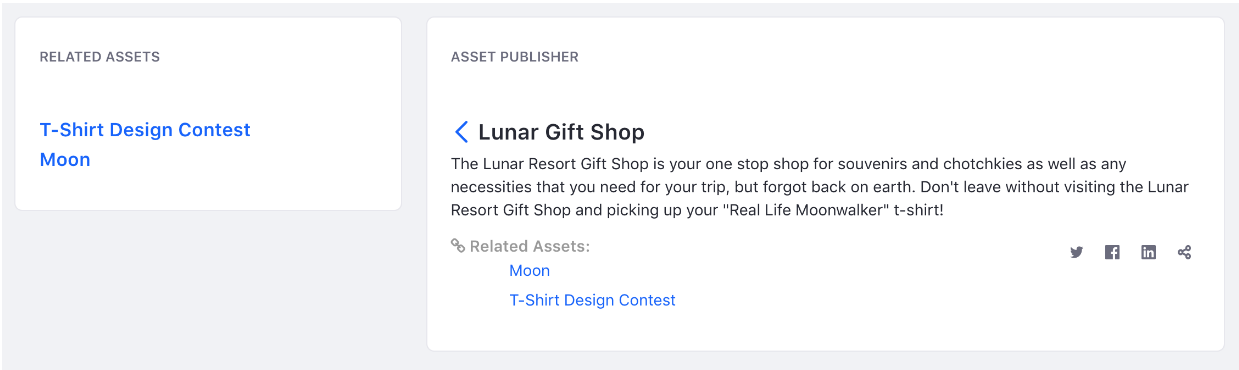

Figure 27.1: Select an asset in the Asset Publisher to see its related assets displayed in the Related Assets application.

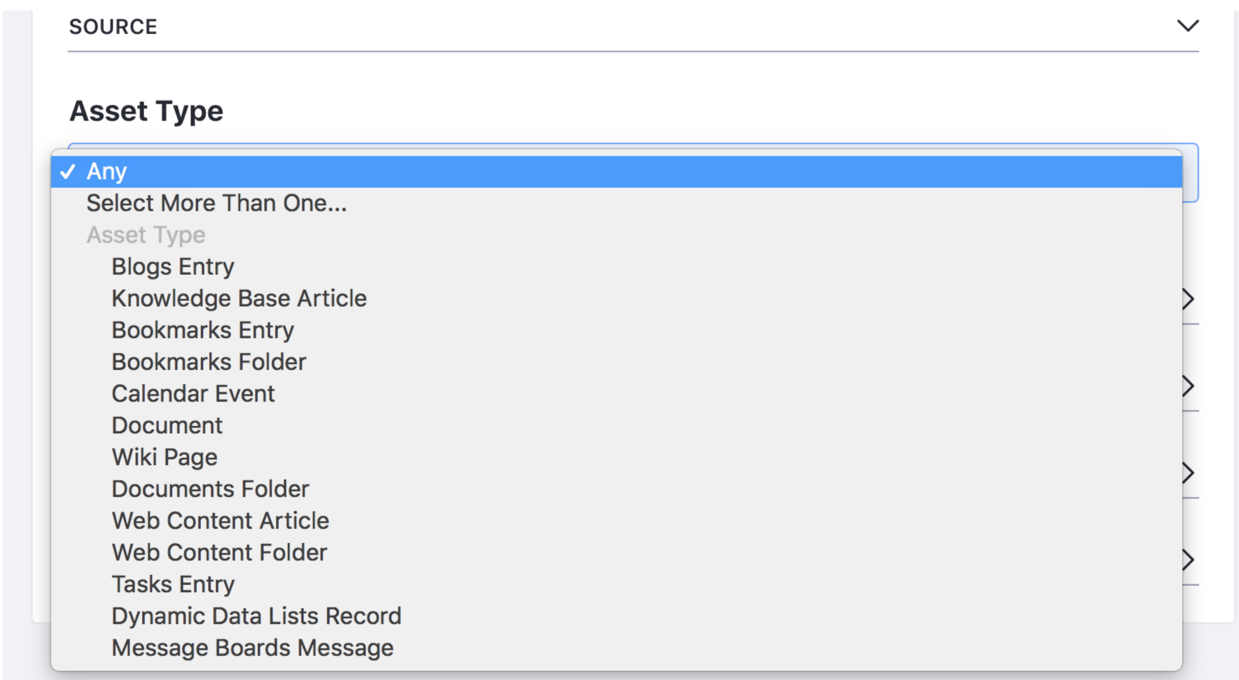

Figure 27.2: Related Assets applications can be configured to display specific content.

Chapter 28

Publishing Assets

As you create web content, remember that pieces of content are assets, just like message board entries and blog posts. Since the Asset Publisher publishes assets, it excels at publishing mixed content types like images, documents, blogs, and of course, web content. This helps in creating a more dynamic web site: you can place user-created wiki entries, blog posts, or message board messages in context with your web content. You'll examine some of the Asset Publisher's features next.

28.1 Querying for Content

The Asset Publisher works by querying for mixed types of content on the fly. Since you can control what and how content is displayed from one location, the Asset Publisher helps to "bubble up" the most relevant content to your users.

To get to all the application's options, click the *Options* icon (⋮) in the application's menu. If you click the *Configuration* option and then *Setup* (if necessary), you can configure the Asset Publisher's settings from the following three areas:

- Asset Selection
- Display Settings
- Subscriptions

Asset Selection configures which assets are displayed. You can set asset selection to *Dynamic* or *Manual*. Dynamic displays assets based on certain rules or filters. For example, you can set the Asset Publisher to display only assets of a certain type or to which certain tags or categories have been applied. Manual asset selection only displays assets that have been explicitly selected by an administrator.

The Asset Publisher supports a scope that restricts both dynamic and manual asset selection. The Asset Publisher can only display assets from its configured scope. By default, the Asset Publisher app is scoped to the site of the page to which it was added. You can, however, customize the scope from the Asset Selection section of the Asset Publisher configuration window. To extend your Asset Publisher's scope,

1. Click *Select* under Scope

2. Choose *Global* to add the global scope, *Pages...* to add the scope to specific pages, or *Other Site...* to add the scope of anther site.

The Display Settings section of the Asset Publisher configuration window is for customizing how content is displayed. The Subscription section enables, disables, or configures email subscriptions and RSS subscriptions. In the following sections, you'll explore the available configurations for the Asset Selection, Display Settings, and Subscriptions sections of the Asset Publisher's configuration window. You'll start by learning how select content manually. You'll see that it's very similar to using the Web Content Display application except that you can select assets of any type, not just web content articles.

28.2 Selecting Assets

You can configure Asset Publisher to select assets manually or dynamically through various criteria. Within those options there is flexibility in what assets are displayed and how they are displayed.

Selecting Assets Manually

To enable manual asset selection,

1. Click the click the *Options* icon ( ) in the application's menu.
2. Select *Configuration* from menu.
3. In the Asset Publisher configuration, select *Manual* from the select box beneath *Asset Selection* sets.

Now you must select a *Scope* and specific *Asset Entries* from that scope to display. You can configure multiple scopes, including the global scope, from which to select assets.

When selecting assets manually, a list of configured scopes appears under the Scope heading. You can configure scope like this:

1. Click the *X* button at the right to remove a scope from the list.
2. Click the *Select* button to add additional scopes to the Asset Publisher's configuration.
3. After you've added a scope, a new Select button appears under the Asset Entries heading. A list of assets selected for display appears in the Asset Entries section. You can select assets to be displayed by clicking the appropriate *Select* button. One button appears for each configured scope. By default, these are the available asset types:

    - Blogs Entry
    - Bookmarks Entry
    - Bookmarks Folder
    - Calendar Event
    - Basic Document
    - Google Docs
    - Contract
    - Marketing Banner

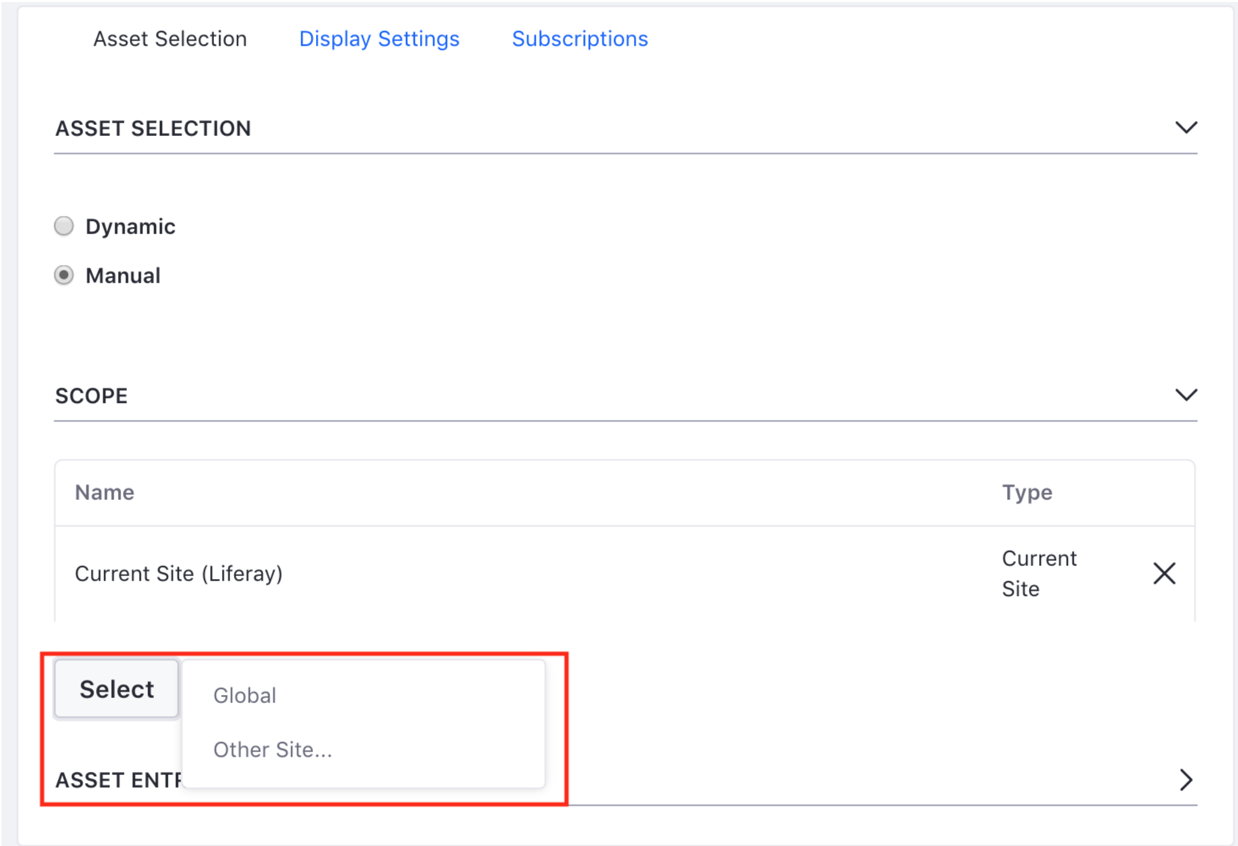

Figure 28.1: Selecting assets in the Asset Publisher manually is similar to selecting assets in the Web Content Display application except that you can select assets of any type, not just web content. You can also add scopes to expand the list of assets that are available to be displayed in the Asset Publisher.

- Online Training
- Sales Presentation
- Documents Folder
- Dynamic Data Lists Record
- Message Boards Message
- Basic Web Content
- Web Content Folder
- Wiki Page

You can select any number of assets to be displayed. Note, however, that there's a display setting called *Number of Items to Display* that determines the maximum number of items to display (or, if pagination is enabled, the maximum number of items to display per page). The Asset Publisher can mix and match different asset types in the same interface.

4. When you're done selecting items to display, click *Save*. Any selected assets are added to the list of assets that are displayed by the application.

Once you have your content selected, you can configure the display types to configure how the content appears. We'll discuss the display settings in more detail after we finish discussing how to select assets for display.

Manual asset selection lets you select assets of various types from different scopes, but it can be time-consuming to update the assets that should be displayed. It's often more convenient to use the Asset Publisher to select content dynamically.

Selecting Assets Dynamically

The Asset Publisher's default behavior is to select assets dynamically according a set of customizable rules. These rules can combined so that they compliment each other to create a nice, refined query for your content. Assets are filtered by permissions automatically, no matter how complicated your asset selection rules are. You have the following rule types:

Scope: Choose the sites containing the content that should be selected. This works the same way as with manual asset selection: assets can only be displayed if they belong to a configured scope. The following scope options are available:

- *Current Site*
- *Global*
- *Other Site*

The Other Site scope option is unavailable for Asset Publisher applications configured on a page template (e.g., Content Display Page).

Asset Type: Choose the asset types you want, from all assets, to only one, or any combination in between. For example, you could choose only web content, only wiki entries, or any combination of multiple types.

Filter: Add as many filters on tags or categories as you like. You can choose whether the content must contain or must not contain any or all of the tags or categories that you enter.

Once you've set up your filter rules for dynamically selecting content, you can decide how the content is displayed.

If you've added custom User profile attributes, you can configure the Asset Publisher to display assets that match them. This setting retrieves assets that have matching categorization. These categories must be from the global context. For example, suppose a User has a custom field called *Location* with the type *Text*. If this attribute is set to *Moon*, you could create a vocabulary called *Location* and a category for the Location vocabulary called *Moon*. Then you could categorize content with *Moon* in the *Location* vocabulary. With this organizational setup, adding an Asset Publisher and specifying *Location* as the Asset Publisher's custom user attribute would only display content that had been categorized as *Moon*. Pretty cool, right?

See Defining Categories for Content for further information.

In addition, you can use these advanced filters:

- **Show only assets with *Welcome* as its display page** displays only assets specifically configured for the *Welcome* page.
- **Include tags specified in the URL?** lets you specify tags in the URL for the Asset Publisher to display.

The *Ordering and Grouping* section of the Asset Publisher precisely controls how content is ordered and grouped when displayed. You can order the assets displayed by Asset Publisher in ascending or descending order by the following attributes:

- Title

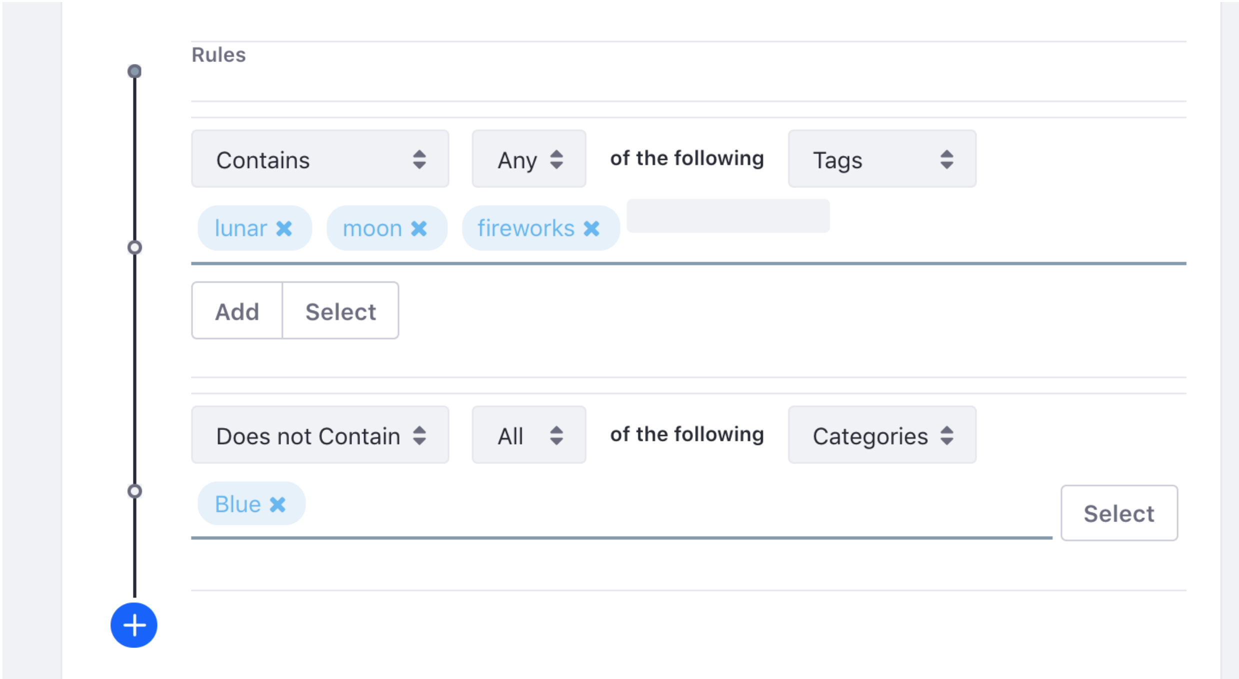

Figure 28.2: You can filter by tags and categories, and you can set up as many filter rules as you need.

- Create Date
- Modified Date
- Publish Date
- Expiration Date
- Priority

Say you have a series of "How To" articles that you want displayed in descending order based on whether the article was tagged with the *hammer* tag. Or suppose you want a series of video captures to appear in ascending order based on a category called *birds*. For these use cases, you can configure the ordering and grouping settings.

You can also configure a second ordering. The second ordering is applied to any assets for which the first ordering wasn't sufficient. For example, if you ordered assets by title and there are multiple assets with the same title, the second ordering takes effect, perhaps the publication date.

You can establish grouping rules as well as ordering rules. You can group assets by type or by vocabulary. For example, suppose you have a vocabulary called *Membership Type* with two categories: *Premium* and *Regular*. If you group assets by Membership Type, all assets with the Premium category appear in one group and all assets with the Regular category appear in another group. Grouping rules are applied before any ordering rules: they're a way to divide up the displayed assets into separate lists. The ordering rules are applied separately to each group of assets.

Note that grouping and ordering rules are only one way to control how your content appears. You can refine the display through many other display settings which you'll examine next.

Note: The following actions have immediate effects in your Asset Publisher: - Change the value of the *Asset Selection* option. - Change the value of the *Scope* option. - Select, add, sort or delete asset entries (only when selecting assets manually).

Other changes happen after clicking *Save*. Next you'll learn about the Asset Publisher's other configuration options.

28.3 Configuring Display Settings

This document has been updated and ported to Liferay Learn and is no longer maintained here.

From the Asset Publisher's configuration page, open the Setup tab's *Display Settings* sub-tab. This section gives you precise control over the display of your assets. There are many options available to configure how you want your content to appear. Many of these, such as printing, flags, ratings, comments, comment ratings, and social bookmarks work the same way they do in the Web Content Display application.

Show Add Content Button: When selected, an *Add New* button appears that lets users add new assets directly from the Asset Publisher application. This is checked by default.

Display Template: This selector lets you choose an application display template to customize how the Asset Publisher displays assets. These templates are in every site by default:

- *Abstracts:* Shows the first 200-500 characters of the content, defined by the **Abstract Length** field. This is the default display template.
- *Table:* Displays the content in an HTML table which can be styled by a theme developer.
- *Title List:* Displays the content's title as defined by the user who entered it.
- *Full Content:* This display template displays the entire content of the entry.

There's also the *Rich Summary* and *Map* display templates that belong to the global scope. The Rich Summary template provides a summary view of each asset along with a *Read More* link to the article's full content. The Map template displays geo-localized assets in either a Google Map or an Open Street Map provider. The map provider can be configured in Instance Settings, and Site Settings in the Advanced section.

Abstract Length: Select the number of characters to display for abstracts. The default is `200`.

Asset Link Behavior: The default value is *Show Full Content*, which displays the full asset in the current Asset Publisher. *View in a Context* causes that asset to be displayed in the application where the asset belongs. For example, a blog entry is displayed in Blogs where it was created. See the section below on display pages for more information.

Tip: When the Asset Publisher displays web content articles with an associated small image, the small image becomes a link to the full article. To use this feature, add or edit a web content article that the Asset Publisher should display. Before clicking *Publish*, click on *Abstracts*, flag *Small Image*, and upload an image. Then click *Publish*. Once your web content article appears in the Asset Publisher's list, clicking the small image takes you to the full article.

Number of Items to Display: Select the maximum number of assets that can be displayed by the Asset Publisher. If pagination is enabled, this number represents the maximum number of assets that can be displayed per page.

Pagination Type: This can be set to *None*, *Simple*, or *Regular*. *None* displays at most the number of assets specified in the **Number of Items to Display** property. *Simple* adds Previous and Next buttons for browsing through pages of assets in the Asset Publisher. *Regular* adds more options and information including First and Last buttons, a dropdown selector for pages, the number of items per page, and the total number of results (assets being displayed).

Show Metadata Descriptions: Enables Metadata descriptions such as *Content Related to...* or *Content with tag...* to be displayed with the published assets.

Show Available Locales: Since content can be localized, you can have different versions of it based on locale. Enabling this option shows the locales available, so users can view the content in their languages.

Set as the Default Asset Publisher for This Page: The Asset Publisher app is an instanceable app: multiple Asset Publishers can be added to a page and each has an independent configuration. The default Asset Publisher for a page is the one used to display web content associated with the page.

Enable ...: Enable/disable the following options for displayed assets:

- Print
- Flags
- Related Assets
- Ratings
- Comments
- Comment Ratings

The Print option adds a *Print* link to the full view of an asset displayed in the Asset Publisher. Clicking *Print* opens a new browser window with a print view of the asset. Enabling flags, related assets, ratings, comments, comment ratings, or social bookmarks add links to the corresponding social features to the view full of the asset in the Asset Publisher.

Tip: An alternate way to add flags, comments, and ratings to a page is through the *Page Flags*, *Page Comments*, and *Page Ratings* applications. Just add the applications in the appropriate location near the asset that should have feedback.

Metadata: Select various metadata types to be displayed (see below). For example, you can select tags and categories for display. Upon saving your configuration, the Asset Publisher displays tags and categories for each displayed asset. Then users can click on the tags and categories to filter the displayed assets manually.

Next you'll learn about configuring subscriptions for email and RSS through the Asset Publisher.

28.4 Configuring Asset Publisher Subscriptions

The Asset Publisher application supports two kinds of subscriptions: RSS subscriptions and email subscriptions. To enable subscriptions, click the Asset Publisher's Options icon and select *Configuration*. In the configuration window, open the Setup tab's Subscriptions tab. There are two options: *Enable RSS Subscription* and *Enable Email Subscription*.

Enabling RSS Subscription creates an RSS feed containing links to all assets that the Asset Publisher is configured to display. A link to this RSS feed appears at the bottom of the Asset Publisher application. This option is only available when the *Dynamic* Asset Selection is configured.

Enabling Email Subscription adds a *Subscribe* link to the Asset Publisher. Users wishing to be notified of newly published assets can click on this link to be added to the subscription list. Liferay DXP periodically checks for new assets and sends emails to subscribed users informing them about the new assets. By default, Liferay performs this check every twenty-four hours.

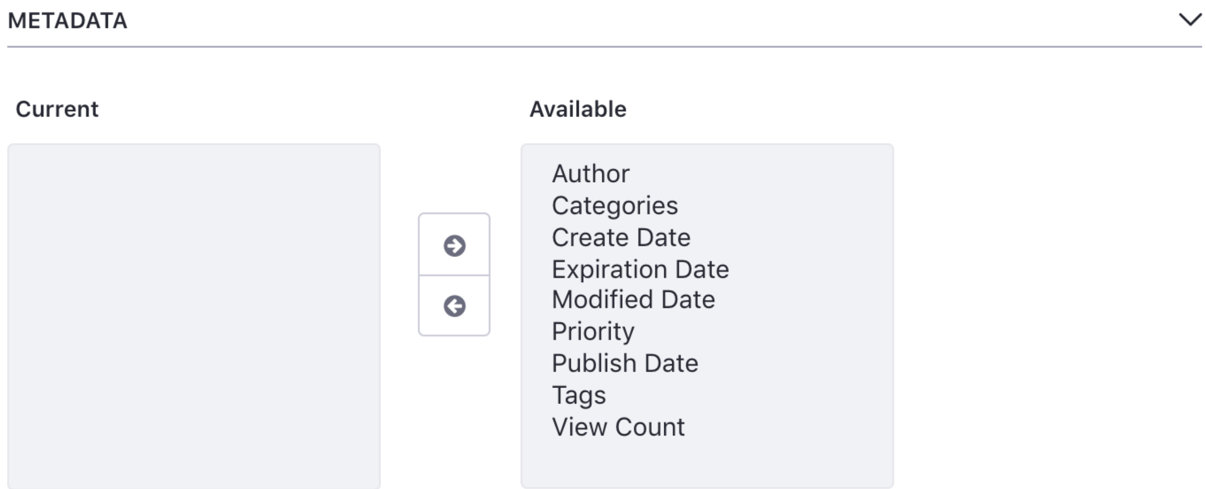

Figure 28.3: You can configure the Asset Publisher to display various kinds of metadata about the displayed assets.

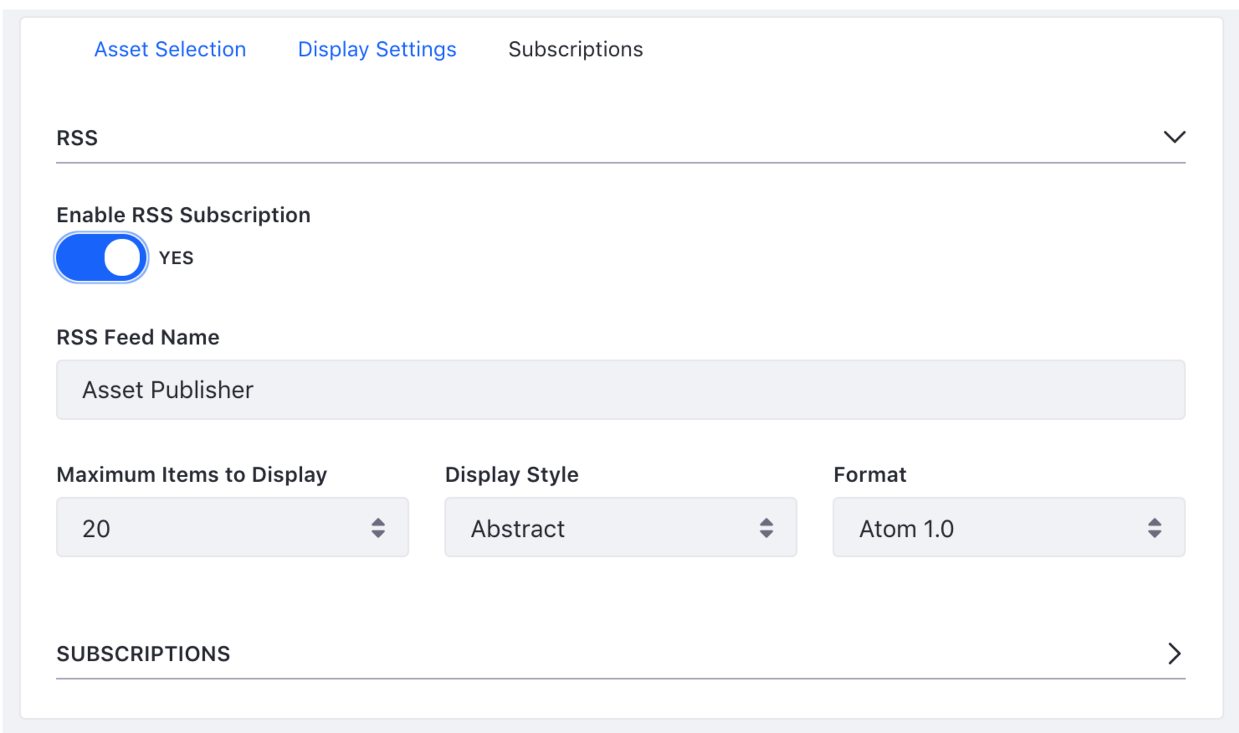

Figure 28.4: When RSS subscriptions have been enabled for an Asset Publisher application, a link to the Asset Publisher's RSS feed appears. Users can subscribe to the Asset Publisher's RSS feed using their preferred RSS reader.

Chapter 29

Publishing RSS Feeds

RSS is a family of web feed formats used to publish frequently updated works such as blog entries and news articles. RSS allows users to stay up-to-date with your site's content without actually having to visit your site. Users can use their own RSS feed readers to aggregate content, and you can also use RSS to share and aggregate content across sites. Next, you'll see how to create RSS feeds from Asset Publisher configurations.

29.1 Configuring RSS Feeds

Note: RSS feeds are deprecated for Liferay DXP 7.1 and are disabled by default. To leverage RSS feeds, you must enable this feature. Go to the Control Panel → *Configuration* → *System Settings* → *Web Content*. Under the *System Scope* → *Administration* tab, check the *Show Feeds* box. For more information on deprecated apps, see this article.

To manage a Liferay site's RSS feeds, navigate to your Site's Site Administration → *Content* page and click *Web Content*. Site administrators can use this Web Content menu option to manage their site's web content, including web content structures and templates, which you learned in the Creating Web Content section. Site administrators can also use this option to manage their site's RSS feeds. To add a new feed:

1. Go to the *Feeds* tab.

2. Click the *Add Feed* button.

3. Enter a *Name*, select a *Target Page*, and select a *Web Content Structure* for the feed.

A feed's target page serves two purposes:

1. The site the target page belongs to determines which web content articles appear in the feed. For example, if the target page belongs to the Marketing site, only web content articles belonging to the Marketing site appear in the feed.

2. The target page is the page where "orphaned" web content articles are displayed. Orphaned web content articles have been published in your Site but are not displayed in specific Web Content Display applications. Liferay RSS feeds can provide links to any published web content articles, both orphaned articles and articles that have been configured to be displayed in specific Web Content Display applications. For articles that have been configured to be displayed, the RSS feeds' links point to the Liferay page of that app. For orphaned articles, the RSS feeds' links point to the feed's target page. When users click on such links for orphaned articles, the full content of the orphaned article is displayed on the target page.

To specify a target page, you must enter the target page's friendly URL. Note that friendly URLs don't include the host name. For example, the friendly URL of a public page called *Welcome* belonging to a Site called *Marketing* might look like this: `/web/marketing/welcome`. Optionally, you can specify a target portlet ID. This would be the portlet ID of a Web Content Display application on the target page in which orphaned web content should be displayed. The application must exist or else the content isn't displayed. The URL field contains the address of your RSS feed. It appears after you've actually created the feed by clicking *Save*.

The final two sections of the *Add Feed* form are for customizing the web content articles that appear in your feed.

1. Web Content Constraints selects a web content structure to filter the articles that appear in your feed. This is useful since all web content articles are created using web content structures.

2. Presentation Settings customizes additional details about your feed and how articles are displayed in your feed. Leave the Feed Item Content set to *Web Content Description* if you want a description of each article to appear in your feed. Set it to *Rendered Web Content: Use Default Template* if you want the full content of each article to appear in the feed. Customizing the Feed Type allows you to choose which web feed language to use for your feed. You can choose *Atom 1.0* (the default), *RSS 1.0*, or *RSS 2.0*. Customize the *Maximum Items to Display* to choose the maximum number of articles should appear in your feed at one time. Leave the Order By Column set to *Modified Date* to have articles arranged in order from the last time they were published or modified. You can set the Order by Column to *Display Date* if you want to have articles arranged in order from the time they were configured to be displayed in a specific Web Content Display application. Lastly, you can leave the Order by Type set to *Ascending* to have the oldest articles at the top of the feed or you can set it to *Descending* to have the newest articles at the top of the feed.

When you're done configuring your RSS feed, click *Save* to create your feed.

Once one or more feeds have been created, they'll appear in a list in the Feeds tab. You can edit existing feeds using the same form used for creating them. The main difference is that when you edit an existing feed, the URL field is populated. Copy this URL into a new browser tab or window to test your feed. From the Feeds popup window, you can also customize the permissions of feeds or delete feeds.

It's possible to completely disable RSS feeds at the instance level. You can do this by setting the `rss.feeds.enabled` property to `false` in your `portal-ext.properties` file. By default, it's set to `true`. If you keep the default, RSS enabled, you can make several other RSS property customizations. Please refer to the RSS section of your `portal.properties` file for details.

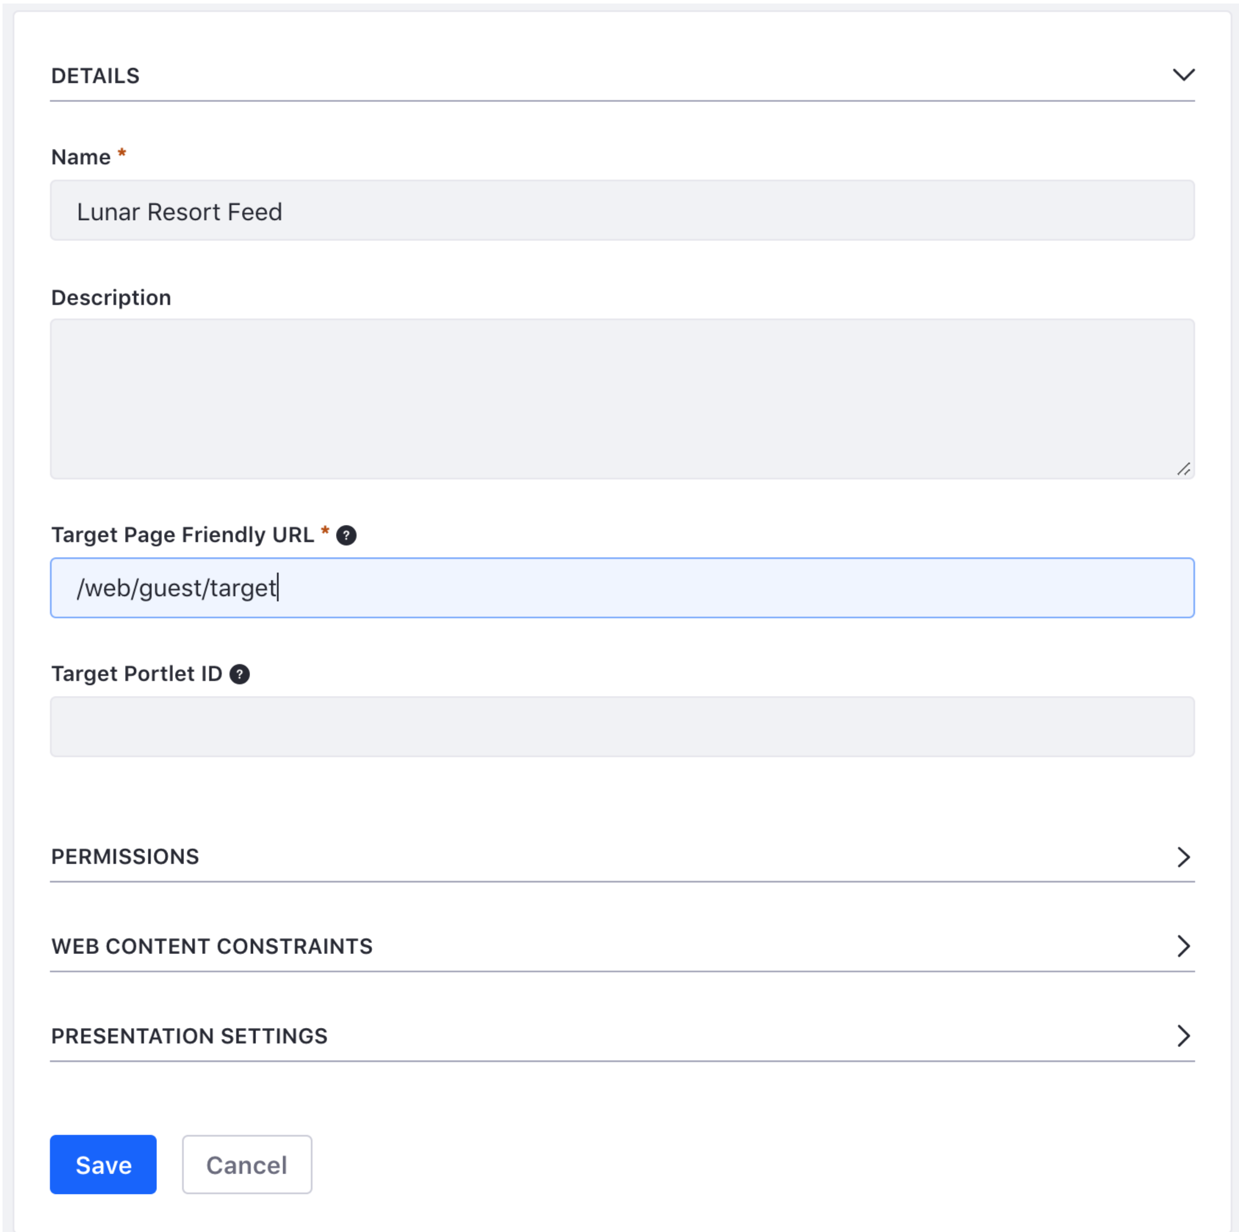

Figure 29.1: To create a new RSS feed, you only need to specify a name, target page, and web content structure. Of course, you can also configure other features of the feed such as its permissions, web content constraints, and presentation settings.

29.2 The RSS Publisher Widget

The RSS Publisher widget displays RSS feeds. If you're looking for a web-based RSS reader, look no further: just add the RSS Publisher widget to one your personal Site's private pages, and *voila*! You have your own personal RSS reader. You can select the RSS feeds the widget displays and how it displays them. The RSS Publisher widget can also be placed on Sites' public or private pages to make feeds available to guests or Site members, respectively. In these cases, make sure that only Site administrators have permission to customize the RSS widget and select feeds to be displayed.

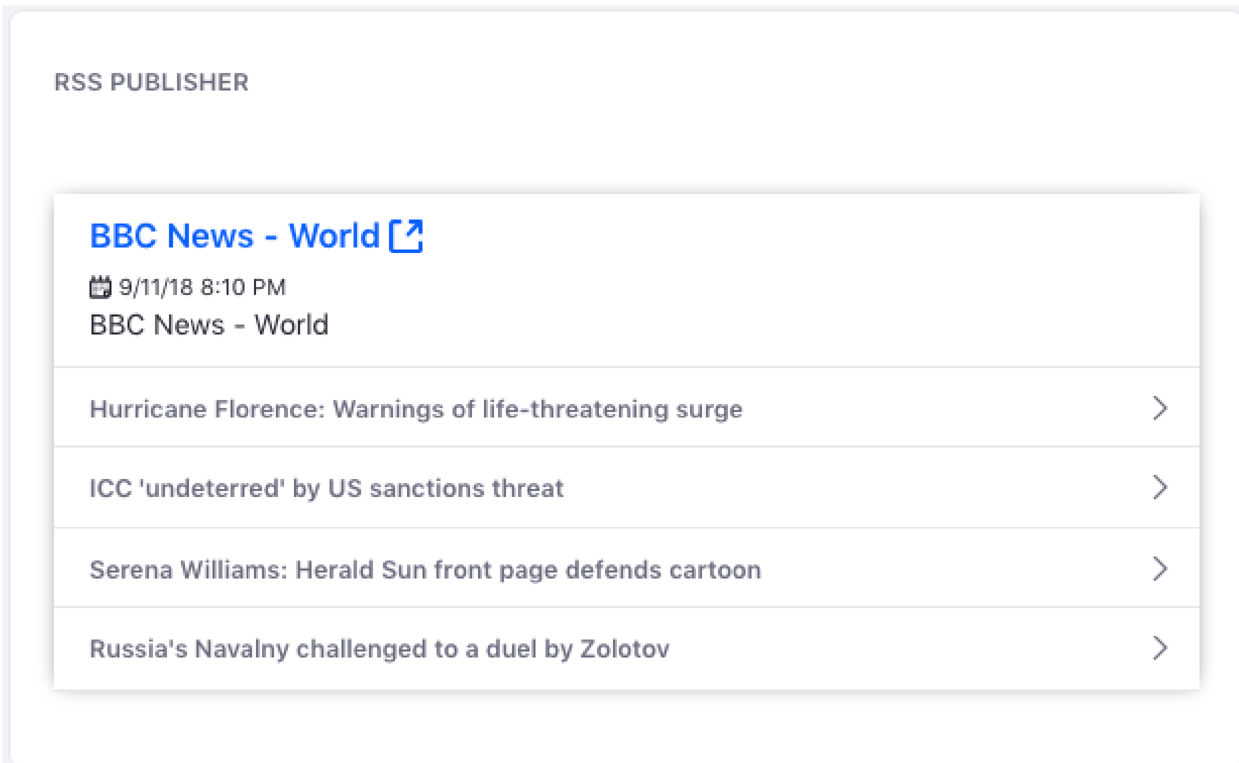

Figure 29.2: The RSS Publisher widget lets you display RSS feeds of your choosing.

Note: If you run your server behind a proxy, you must set the appropriate Java proxy settings (such as `http.proxyHost=` and `http.proxyPort=`) in your `setenv` script or in your `system-ext.properties`. Without these properties, the RSS Publisher widget can't access any RSS feeds.

Note that the RSS Publisher widget is deprecated. In Liferay CE Portal 7.1 GA2+, and Liferay DXP 7.1 FP4+, the widget is available from the *Add* (➕) → *Widgets* → *News* menu. However, the widget is hidden in earlier releases of Liferay CE Portal 7.1 and Liferay DXP 7.1. In these releases, you must therefore make the widget visible via a configuration file. The next section shows you how to do this.

Ensuring that the RSS Publisher Widget is Visible

The RSS Publisher widget is hidden in Liferay CE Portal 7.1 GA1 and versions of Liferay DXP 7.1 prior to Fix Pack 4. If you're running one of these versions, follow these steps to ensure that the widget is visible:

1. Shut down your portal instance, if it's running.
2. In `[Liferay Home]/osgi/configs`, create the file `com.liferay.rss.web.internal.configuration.RSSWebCacheConfig` Note that the Liferay Home folder is usually the application server's parent folder.
3. Add the following configuration to the file you created in the first step:

   ```
   com.liferay.portlet.display-category="category.news"
   ```

4. Start Liferay DXP.

The RSS Publisher widget is now available in the *Add* () → *Widgets* → *News* menu.

Using the RSS Publisher Widget

You can add the RSS Publisher widget to a page from the *Add* () → *Widgets* → *News* menu. Once you've done so, open the widget's Configuration menu by clicking on the *Options* icon () at the top-right corner of the widget and selecting *Configuration*.

By default, the RSS Publisher widget displays one feed. In the *Feeds* section, add or remove a feed via the plus or minus buttons, respectively. To add a feed, enter its URL and title in the respective fields. If you leave the *Title* field blank, the feed's default title is used (the *Title* field is for custom titles).

In the top section, use the following toggles to enable/disable the display of the feed's details:

- Show Feed Title
- Show Feed Published Date
- Show Feed Description
- Show Feed Image
- Show Feed Item Author

You can also select the number of entries and expanded entries that should be displayed per feed. Expanded entries show more of an article's actual content than regular entries. By default, each feed shows four entries per feed and eight expanded entries per feed. You can set the feed image alignment to control whether feed images appear to the right or left of the text. By default, the feed image alignment is set to *Right*.

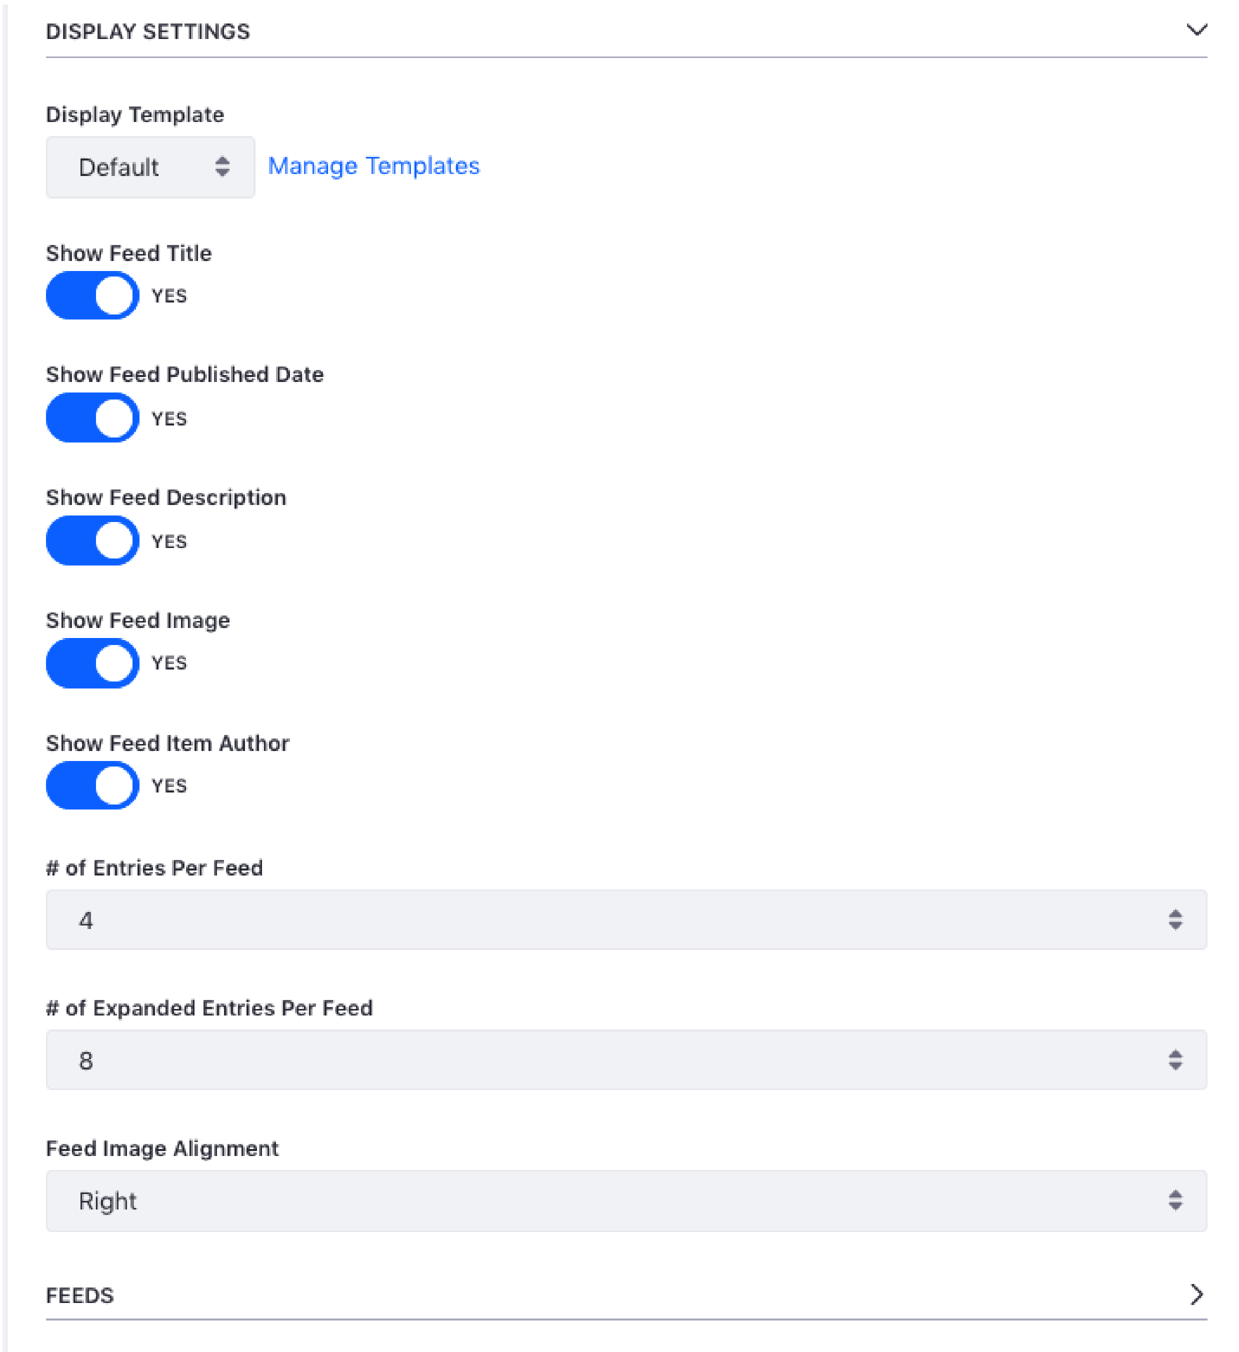

Figure 29.3: The RSS Publisher widget's configuration lets you customize how the widget displays RSS feeds.

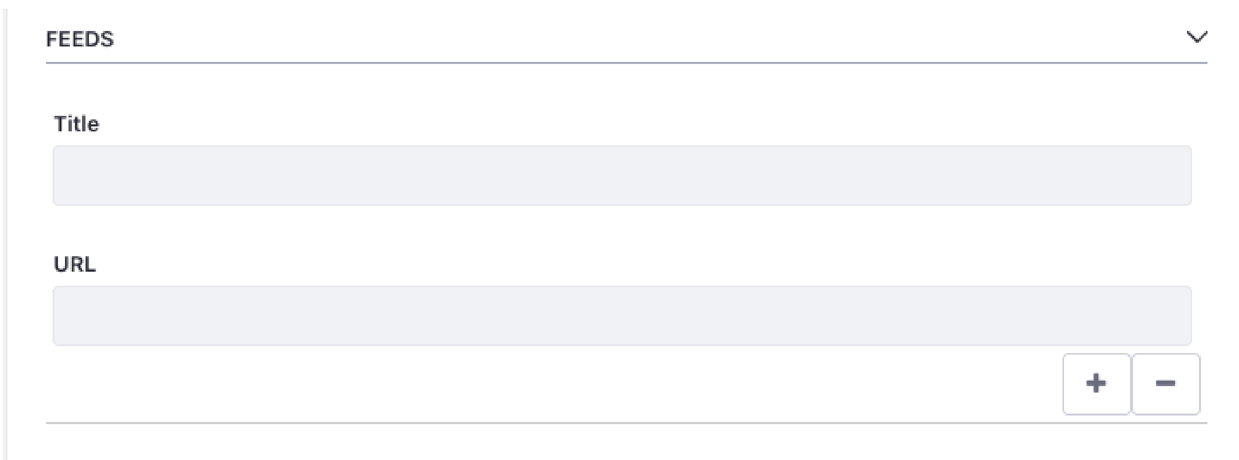

Figure 29.4: You can also use the RSS Publisher widget's configuration to specify which feeds to display.

Chapter 30

Configuring Widgets

Just like siblings have common features inherited from their parents, widgets that ship with Liferay DXP also share common features. These include look and feel, exporting/importing app data, communication, sharing, permissions, scoping, and configuration templates. These features work together to facilitate information flow within Liferay DXP and provide an enhanced experience for your users. You'll start with look and feel configuration options.

30.1 Look and Feel Configuration

To access the look and feel configuration menu of any widget,

1. Click *Options* (⋮) in the top right corner of the widget.
2. Select *Look and Feel Configuration*.

Look and Feel Configuration has six tabs:

- General
- Text Styles
- Background Styles
- Border Styles
- Margin and Padding
- Advanced Styling

After making customizations, click *Save* and refresh your page to apply your changes. If you don't like the effect of your changes, some tabs have a *Reset* button to discard changes.

General Settings

On the General tab are the following options:

Use Custom Title enables changes to your widget's title. The value in the title box is displayed on widget's decorator. The title is localizable, so you can provide translations of the title for different languages.

Application Decorators gives you the choice between three decorators: *Barebone*, *Borderless*, and *Decorate*. The Decorate application decorator is the default. Be careful about turning widget borders off; some themes assume widget borders are turned on and may not display correctly with them turned off.

Note: The *Link Portlet URLs to Page* option has been deprecated in 7.0. It can be enabled through system settings in *Control Panel* → *Configuration* → *System Settings* → *Widget Tools* → Look and Feel Configuration*.

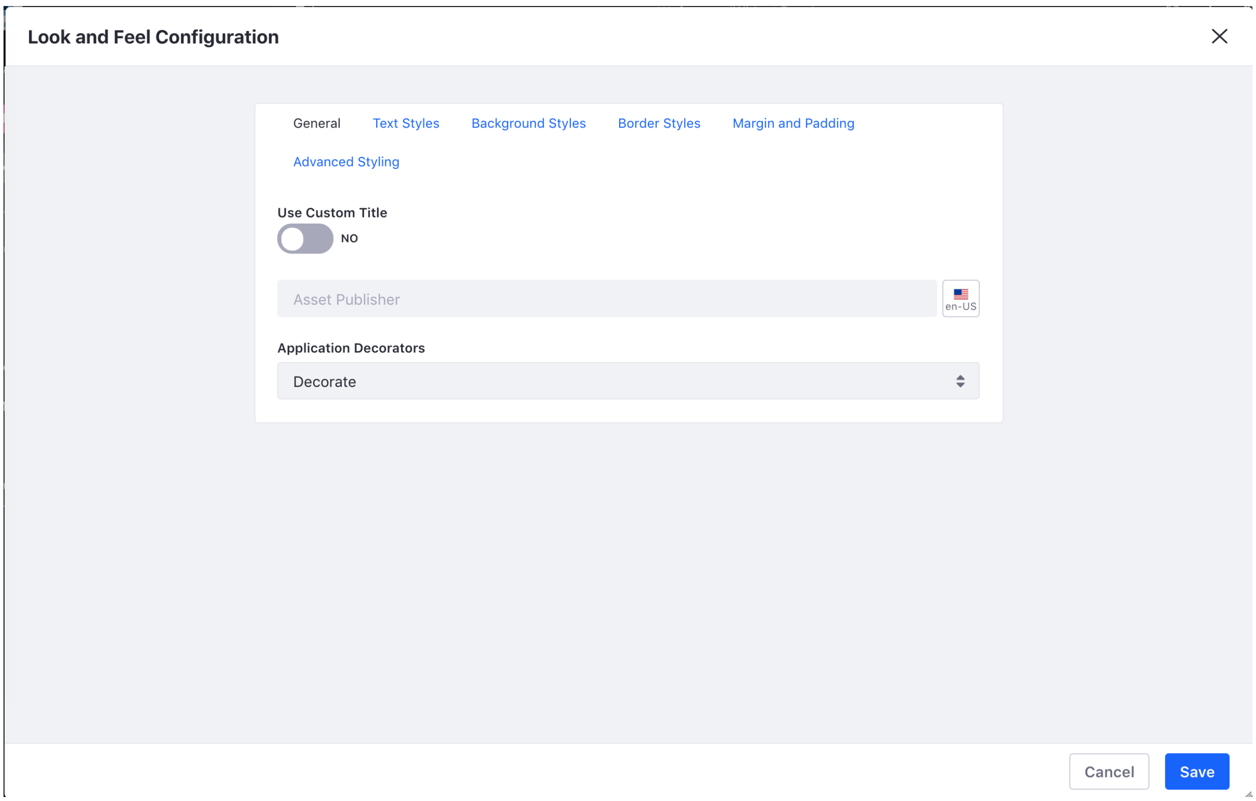

Figure 30.1: The General tab of the Look and Feel Configuration menu allows you to define a custom widget title, link widget URLs to a specific page, and select the widget contrast option using decorators.

Text Styles

Text Styles configures the format of the text that appears in the widget. The options include

Font: Choose various fonts. You can set the text to bold, italics, or both.

Size: Set the font size anywhere from 0.1 em to 12 em, with 0.1 em increments. 1 em is the default.

Color: Set to any six digit hex color code. Click on the text box to open the color palette.

Alignment: Set to *Left*, *Center*, *Right*, or *Justified*.

Text Decoration: Set to *Underline*, *Overline*, or *Strikethrough*. The default text decoration is *None*.

Word Spacing: Set from -1 em to 0.95 em, with 0.05 em increments. 0 em is the default.

Line Spacing: Set from 0 em to 12 em, with 0.1 em increments. 0 em is the default.

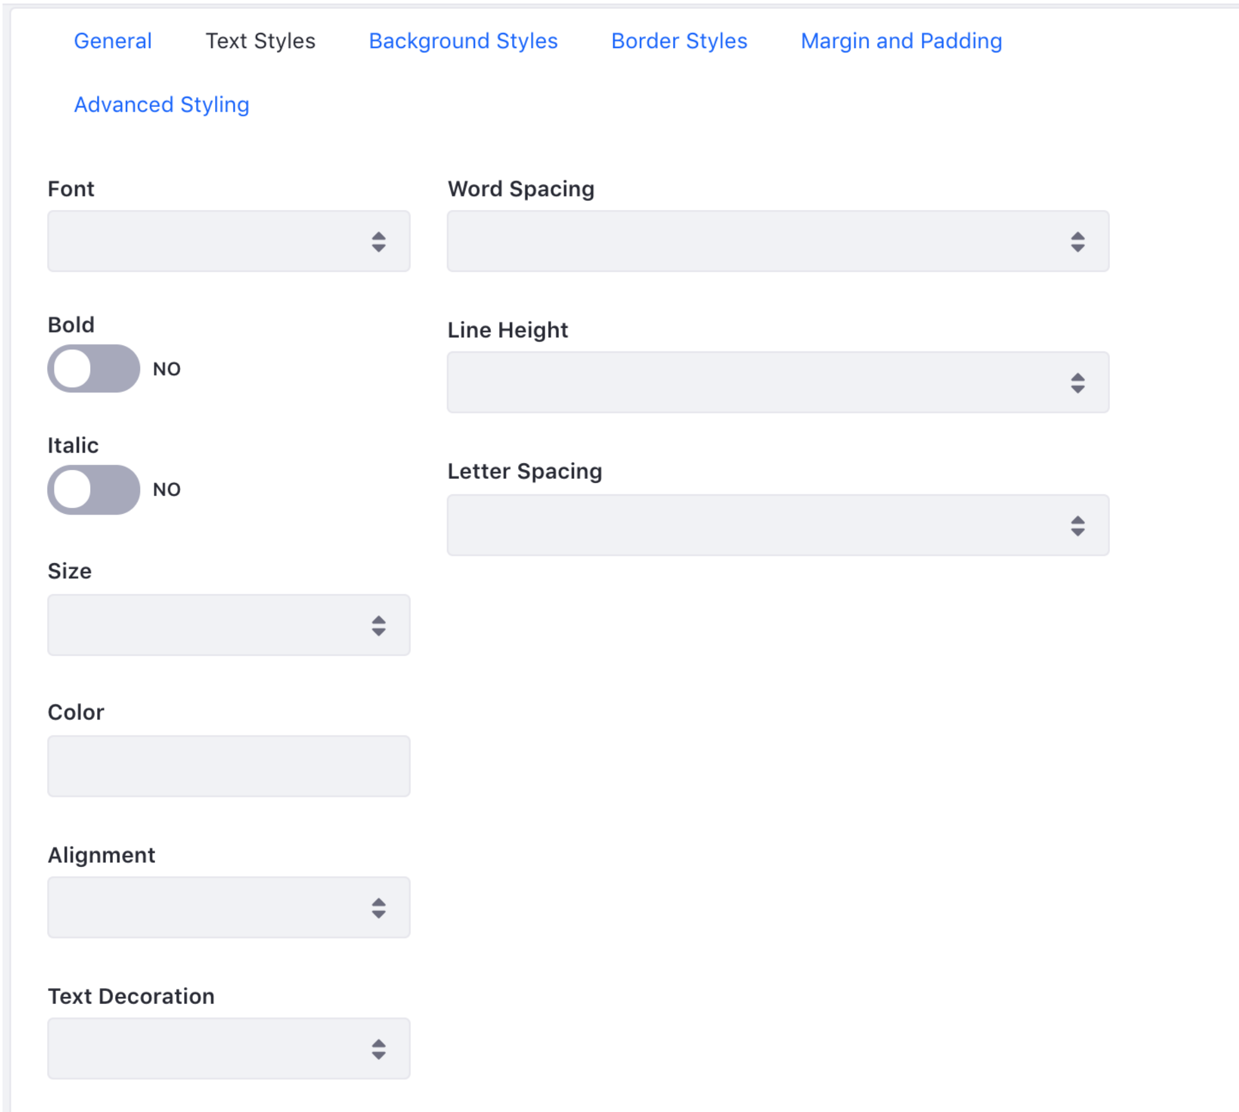

Figure 30.2: The Text Styles tab lets you configure the format of the text that appears in the widget.

Letter Spacing: Set from -10 px to 50 px, with 1 px increments. 0 px is the default.

Background Styles

The Background Styles tab specifies the widget's background color. When you select the text space, you're given a color palette to choose your background color or you can manually enter any six digit hex color code.

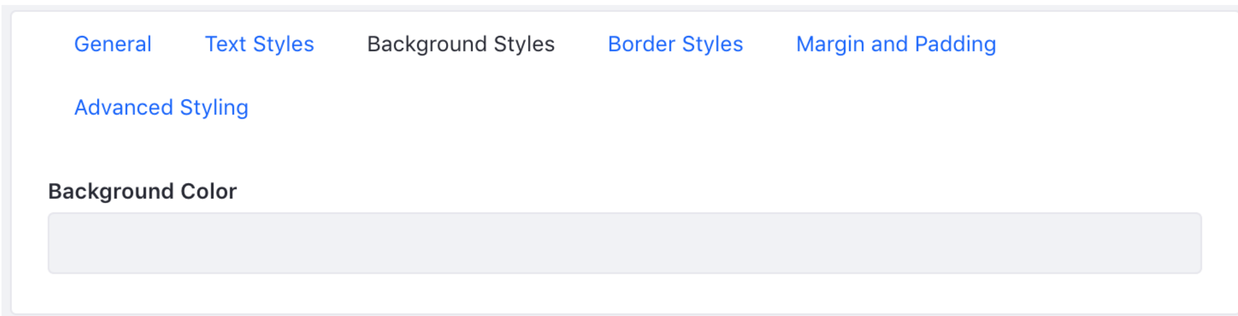

Figure 30.3: The Background Styles tab lets you specify the widget's background color.

Border Styles

The Border Styles tab, configures your widget's border width, style, and color. For each of these attributes, leave the *Same for All* selector enabled to apply the same settings to top, right, bottom, and left borders.

For border width, you can specify any % value, em value, or px value. For border style, you can select Dashed, Double, Dotted, Groove, Hidden, Inset, Outset, Ridge, or Solid. For border color, you can enter any six digit hex color code, just like for the text color and background color. You can also use the color palette.

Margin and Padding

The Margin and Padding tab specifies margin and padding lengths for the edges of your widget. Just like for border styles, leave the *Same for All* selector enabled to apply the same settings to each side (top, right, bottom, and left) of the widget.

For both padding and margin, you can specify any % value, em value, or px value.

Advanced Styling

The Advanced Styling tab displays current information about your widget, including your widget's Liferay ID and CSS classes.

You can also enter custom CSS class names for your widget and custom CSS code. Clicking the *Add a CSS rule for just this portlet* or *Add a CSS rule for all portlets like this one* links adds the CSS code shells into your custom CSS text box. If you check the *Update my styles as I type* box, your CSS code is applied dynamically to your widget so you can see the effects of your edits.

Next, you'll learn about communication between widgets.

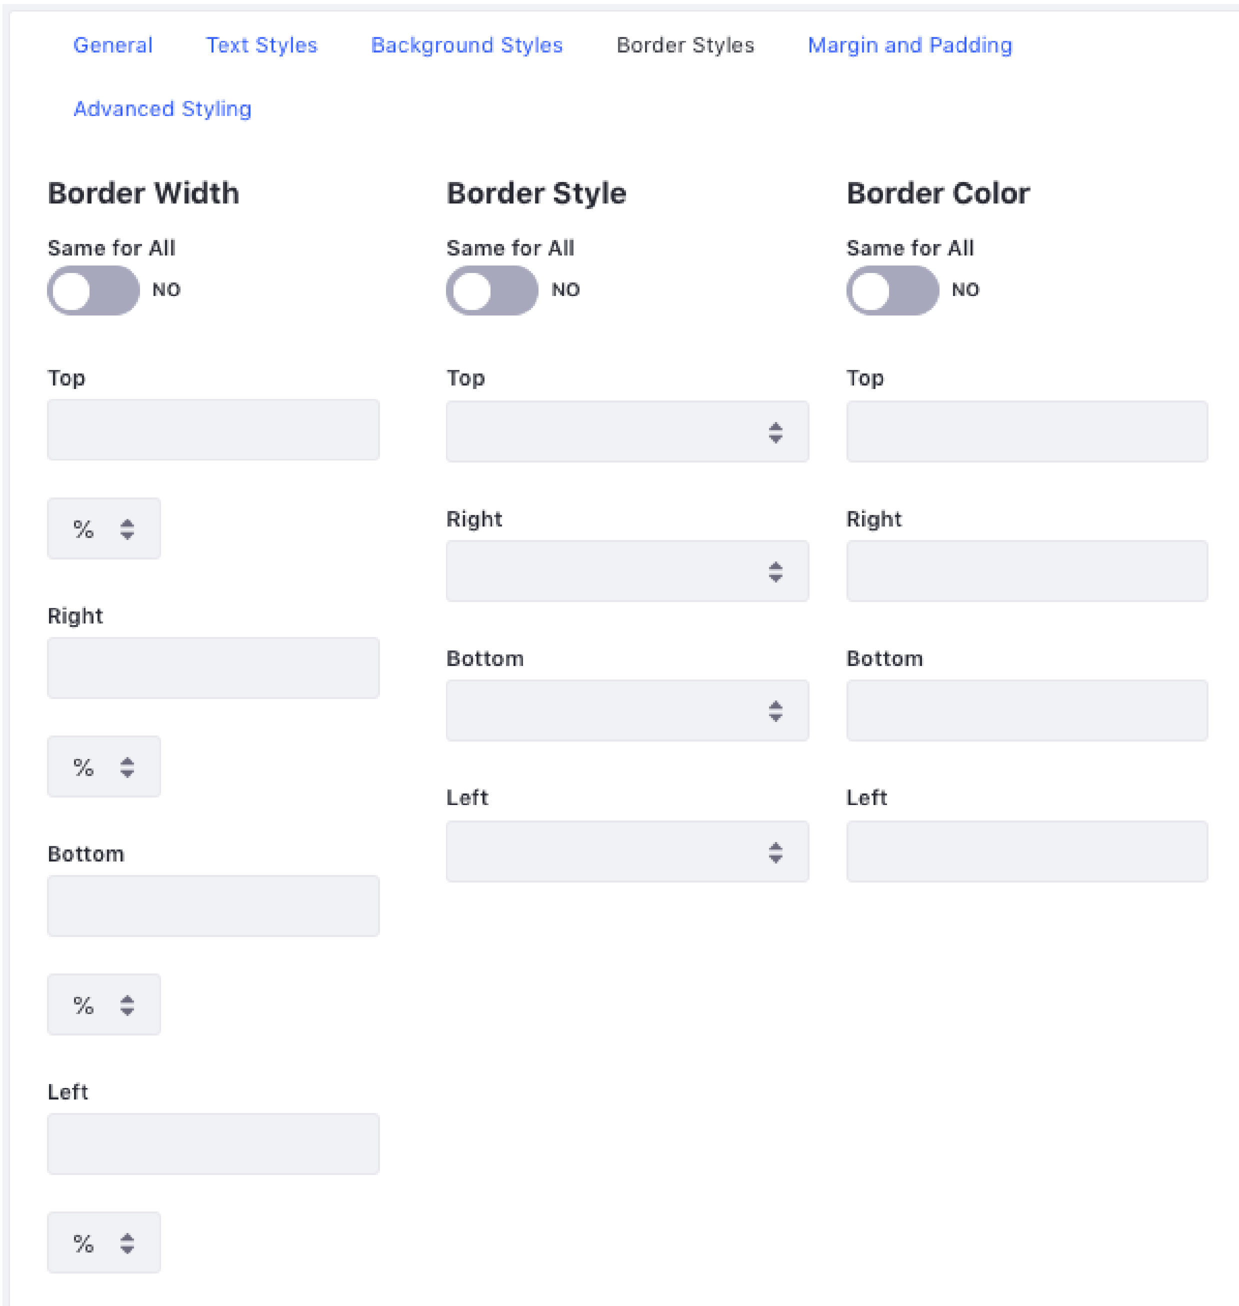

Figure 30.4: The Border Styles tab lets you specify a border width, style, and color for each side of the widget.

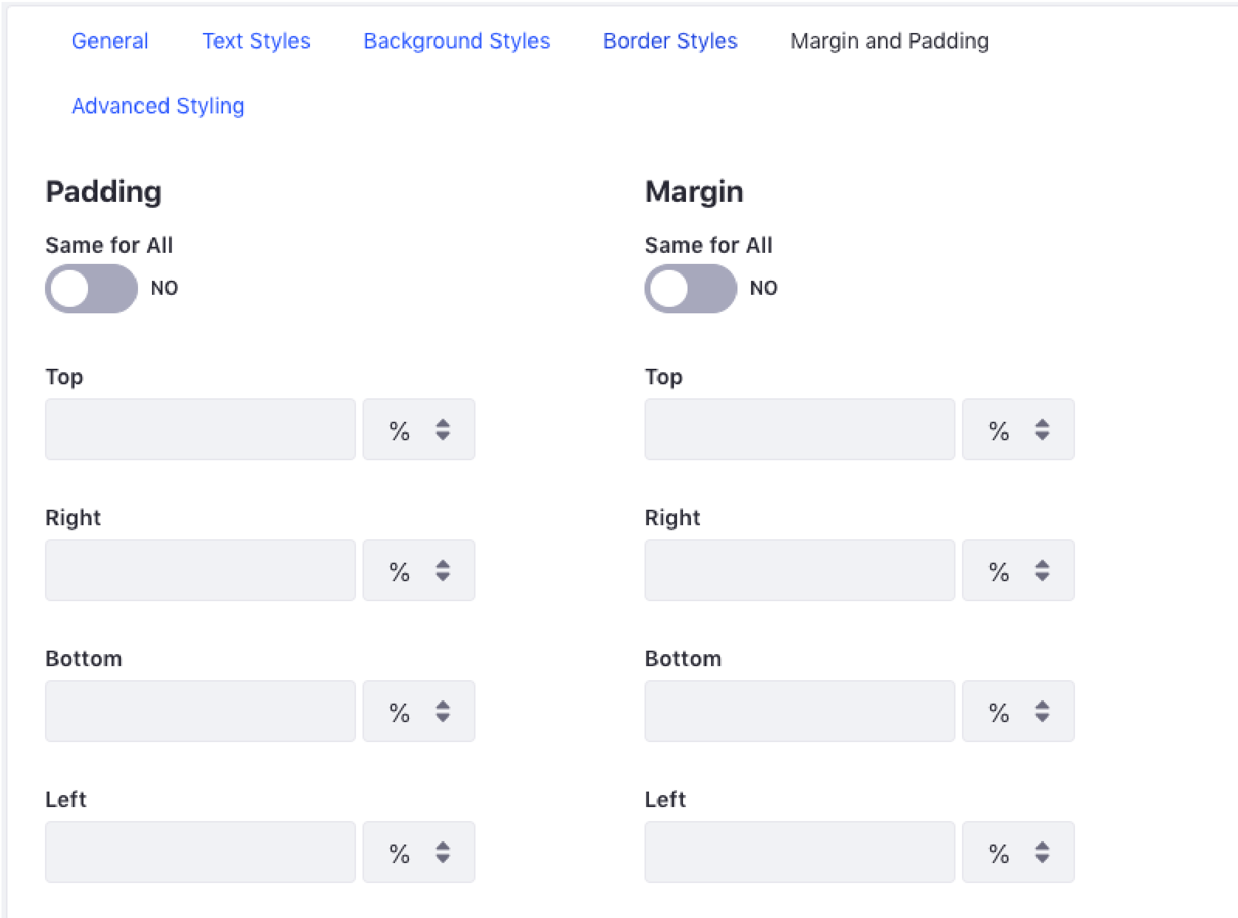

Figure 30.5: The Margin and Padding tab allows you to specify margin and padding lengths for the sides of your widget.

30.2 Exporting/Importing Widget Data

You may need to export data from a specific widget instance, without regard to content on the rest of the Site. There are many widgets that let you export or import their data individually:

- Blogs
- Bookmarks
- Dynamic Data Lists
- Forms
- Knowledge Base
- Message Boards
- Web Content
- Wiki
- And more

Exporting widget data produces a `.lar` file that you can save and import into another widget of the same type. To import widget data, you must select a `.lar` file. Be careful not to confuse widget-

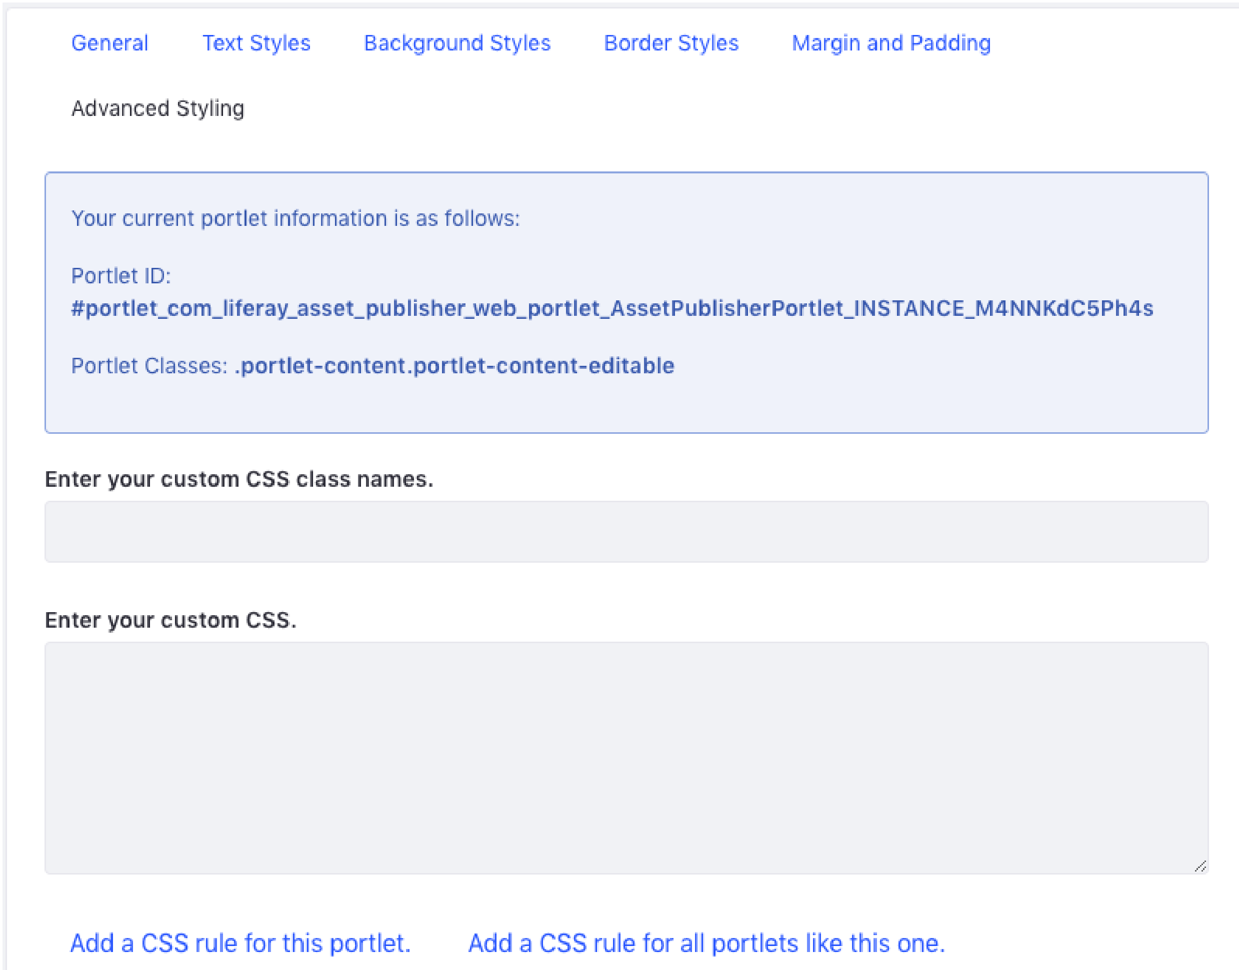

Figure 30.6: The Advanced Styling tab displays your widget's Liferay ID and allows you to enter CSS code to customize the look and feel of your widget.

specific `.lar` files with Site-specific `.lar` files. See the Importing/Exporting Pages and Content article for information on importing/exporting Site page data.

There are two ways to export/import widget content. You can navigate to the widget's administrative area located in the Product Menu, or you can visit the widget on its page. Both export/import menus work the same, but the administrative area may hold content different from its widget counterpart (e.g., Web Content Admin in Product Menu and Web Content Display widget do not offer same content for export/import), so be wary of your selection.

To export or import data from the widget's administrative area, follow the steps below.

1. Navigate to the widget's designated area in the Product Menu. For example, if you plan to export Web Content data, navigate to *Content* → *Web Content*.

2. Click the *Options* button (⋮) from the top right of the page and select *Export/Import*.

3. Select the *Export* or *Import* tab to begin configuring the respective process.

To export or import data from a widget, follow the steps below:

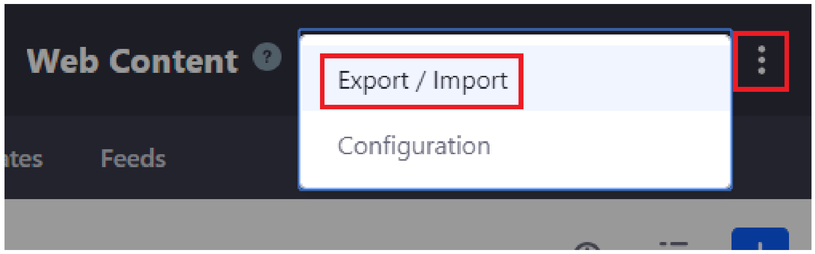

Figure 30.7: You can access a widget's administrative *Export/Import* feature by selecting its Options menu.

1. Ensure the widget you're exporting/importing from is available on a page. You can add widgets from the *Add* (⊞) → *Widgets* menu.

2. Select the widget's *Options* button (⋮) and select *Export/Import*.

3. Select the *Export* or *Import* tab to begin configuring the respective process.

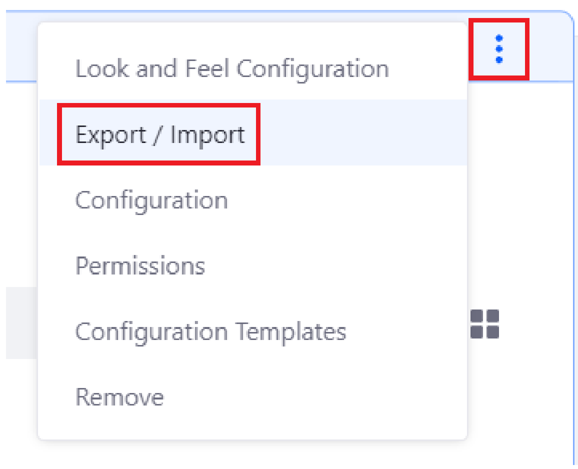

Figure 30.8: You can access a widget's *Export/Import* feature by selecting its Options menu.

Now that you know how to navigate to the *Export/Import* menus, you can explore the export process.

Exporting Widget Data

To export widget data, create a new export process by selecting the *New Export Process* tab (default). You have several export options to configure.

First, you can choose to export your widget's configuration settings. This exports your customized settings from your widget's *Options* → *Configuration* menu. For some widgets, the configuration export might also include content. For example, a Web Content Display widget that shows a web content article also exports the article when exported, even though no content is selected. This applies when publishing a Web Content Display widget too; the configured article is published with the widget.

Next, you can select a *Date Range* of content that you want to export. Content added to your widget within your specified date range is included in the `.lar` file. The following date range choices are available:

All: Publishes all content regardless of its creation or last modification date.

Date Range: Publishes content based on a specified date range. You can set a start and end date/time window. The content created or modified within that window of time is published.

Last...: Publishes content based on a set amount of time since the current time. For example, you can set the date range to the past 48 hours, starting from the current time.

By checking the *Content* box, you can choose specific content you want to export. When you check the *Content* box, more options appear, letting you choose specific kinds of metadata to include. For example, if you have a wiki page with referenced content that you don't want, check the *Wiki Pages* checkbox and uncheck the *Referenced Content* checkbox. Another option is the selection of content types. Two familiar content types in your Liferay instance are *Comments* and *Ratings*. If you want to include these entities in your `.lar` file, select *Change* and select them from the checklist. For more information on managing content types, see the Managing Content Types in Staging article.

Next, you can choose to export individual deletions. This lets delete operations performed for content types be exported to the LAR file.

Finally, you can choose whether to include permissions for your exported content. The permissions assigned for the exported widget window are included if you enable the *Export Permissions* selector.

After you've exported your widget's data, switch to the *Current and Previous* tab to view ongoing export processes and the history of past exports. You can also download the exported `.lar` file from this tab.

Importing Widget Data

To import widget data, you can select the LAR using your file explorer or by dragging and dropping the file between the dotted lines.

Your LAR file is uploaded and displayed to you for review. Click *Continue*.

Now that you've uploaded and confirmed your LAR file, you're given a similar screen to what you'd be offered during export. Several of these options are covered in great detail in the Importing/Exporting Pages and Content tutorial. There are some additional options available: *Update Data* and *Authorship of the Content*. Here's options and descriptions for each section:

Update Data

Mirror: All data and content inside the imported LAR is newly created the first time while maintaining a reference to the source. Subsequent imports from the same source updates entries instead of creating new entries.

Mirror with overwriting: Same as the mirror strategy, but if a document or an image with the same name is found, it is overwritten.

Copy as New: All data and content inside the imported LAR is created as new entries within the current Site every time the LAR is imported.

Authorship of the Content

Use the Original Author: Keep authorship of imported content whenever possible. Use the current user as author if the original one is not found.

Use the Current User as Author: Assign the current user as the author of all imported content.

Once you've selected the appropriate options, select *Import* and your widget's data is imported and ready for use.

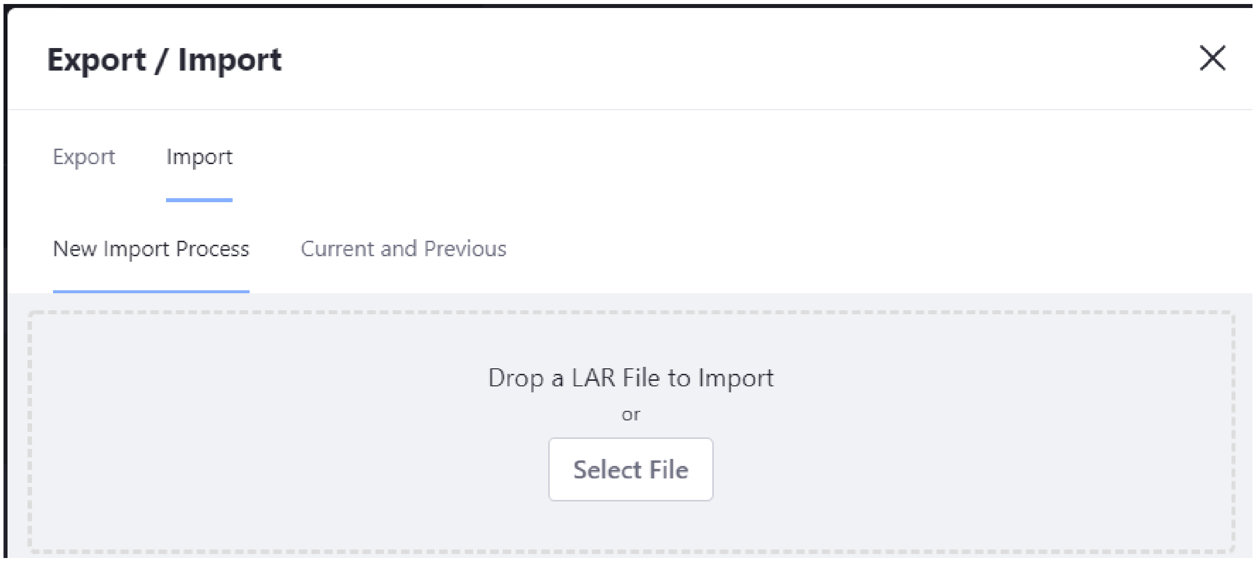

Figure 30.9: When importing widget data, you can choose a LAR file using the file explorer or drag and drop the file between the dotted lines.

30.3 Communication Between Portlet Widgets

Portlet widgets can communicate with each other using public render parameters and events. Public render parameters are easy to use and can be quite powerful. Some Liferay portlets provide a configuration UI to help you get the most out of this communication mechanism. To access this UI, open your portlet's configuration window by clicking on the *Options* icon () and selecting *Configuration*. Then click on the *Communication* tab.

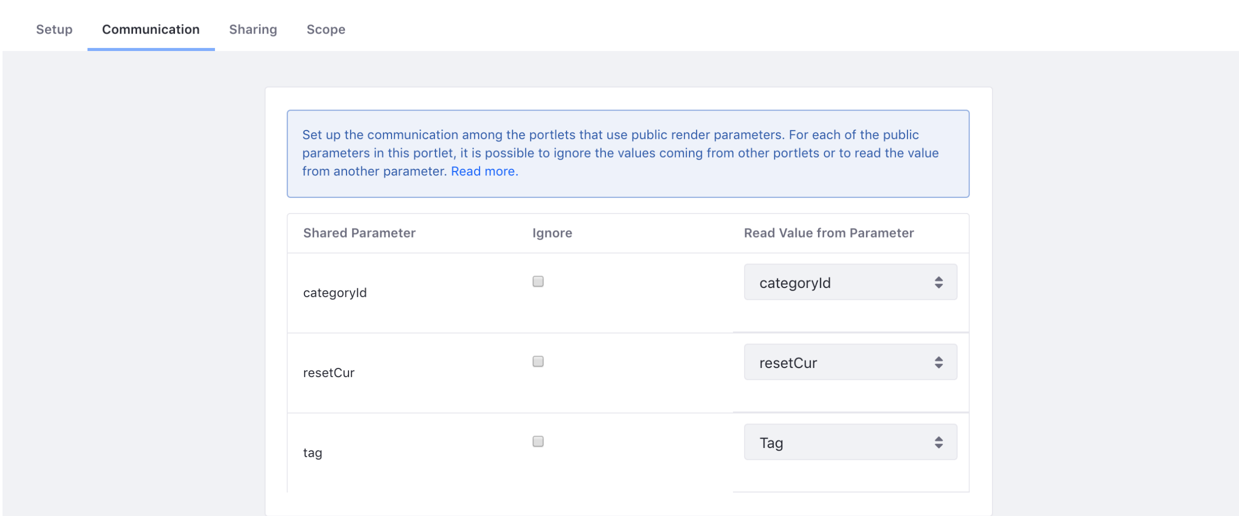

Figure 30.10: You can configure portlets to communicate with each other using public render parameters.

Note: If your widget isn't a portlet, this feature isn't available.

The screenshot above is for the Wiki, which has six public render parameters: `categoryId`, `nodeId`, `nodeName`, `resetCur`, `tag`, and `title`. For each of these parameters, you can configure the portlet to ignore the values coming from other portlets or read the value from another parameter.

Why might it be useful to ignore the values for certain parameters that come from other portlets? Consider a common use case for the Wiki application. The Wiki portlet is often used along with the Tags Navigation portlet so that when a user clicks on a tag of the latter, the Wiki shows a list of pages with that tag. An administrator may want the Wiki to show the front page always independently of any tag navigation done through other portlets. Ignoring the values of the parameter coming from other widgets let this happen.

Reading the value of a parameter from another portlet is an advanced but very powerful option that allows portlets to communicate with each other even if their developers didn't intend them to. For example, imagine that the Wiki is used to publish information about certain countries, and there's another portlet that allows browsing countries for administrative reasons. The second portlet has a public render parameter called *country* with the name of the country. You'd like the Wiki to show the information from the country that's selected in the administration portlet. This can be achieved by setting the value of the title parameter of the Wiki portlet to be read from the country parameter of the administration portlet. Cool, isn't it?

30.4 Sharing Widgets with Other Sites

You can share widgets with other Sites by embedding an instance of a widget running on your Site into another website, such as Facebook. This opens up a whole new avenue of exposure to your web site that you would not have had otherwise. In fact, this is how all those Facebook games work.

To share one of your widgets, open the *Configuration* dialog box from the widget's *Options* icon () and select the *Sharing* tab. There are five sub-tabs under sharing: Any Website, Facebook, OpenSocial Gadget, and Netvibes.

Any Web Site

Copy and paste the provided snippet of JavaScript code into the web site where you want to add the widget. That's all you need to do. When a user loads the page on the other website, the code pulls the relevant widget from your Site and displays it.

Facebook

You can add any widget as a Facebook app. To do this, you must first get a developer key. A link for doing this is provided to you in the Facebook tab. You must create the application on Facebook and get the API key and canvas page URL from Facebook. Once you've done this, you can copy and paste their values into the Facebook tab. Save the configuration and navigate back to the Facebook tab in Liferay DXP. You're given the Callback URL, which you can copy and paste into Facebook. When opening your app in Facebook, the correct callback URL is used to render the application. You can also enable the *Allow users to add [application-name] to Facebook*. Then you can navigate to your app's Options menu and select *Add to Facebook*.

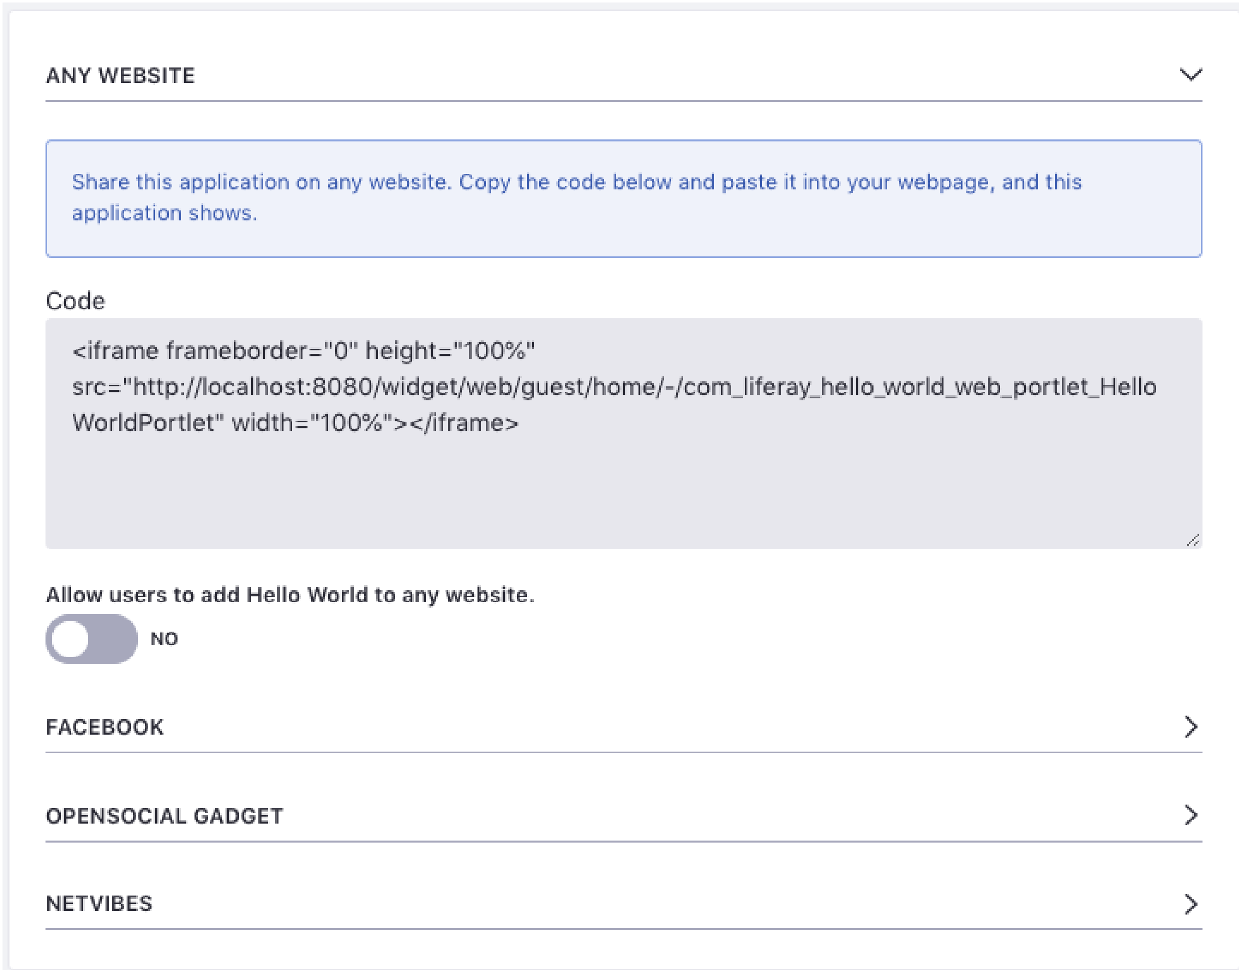

Figure 30.11: The Sharing tab in your widget's Configuration menu lets you share your widget in a variety of ways.

OpenSocial Gadget

OpenSocial comprises a container and a set of APIs for social networking and other web applications. Liferay DXP can serve up applications to be used as OpenSocial Gadgets on any OpenSocial-compatible pages.

To serve a Liferay widget on an OpenSocial platform, copy and paste the provided gadget URL and add it to the appropriate configuration page of the OpenSocial platform you're using. Your Liferay instance serves that widget directly onto that platform's page. The URL provided is unique to the specific instance of the widget, so you could serve multiple instances of the same widget as different OpenSocial Gadgets.

From the Sharing tab in the Configuration menu, you can also enable the selector *Allow users to add [application-name] to an OpenSocial platform*. Click *Save* and revisit the *Options* button of your widget. A new button appears named *Add to an OpenSocial Platform*. When selecting this new button, the URL is provided for sharing the widget to an OpenSocial platform.

Netvibes

Netvibes offers a similar environment where users can log in, create their own personal dashboard, and add customizable widgets to it. To set up Netvibes support for a widget, enable the *Allow users to add [application-name] to Netvibes pages* selector. You can then use the provided URL to create a custom Netvibes widget based on the instance of the Liferay widget that you're using.

Next, you'll learn how to set permissions for Liferay applications.

30.5 Widget Permissions

All of Liferay's widgets support Liferay DXP's robust, fine-grained permissions system. Some higher level permissions can be configured in the permissions tab of the widget's configuration dialog box. You can grant Roles permission to

- Add a display template
- Add the widget to a page
- Configure the app
- Modify the widget's permissions
- Modify the widget's preferences
- View the widget

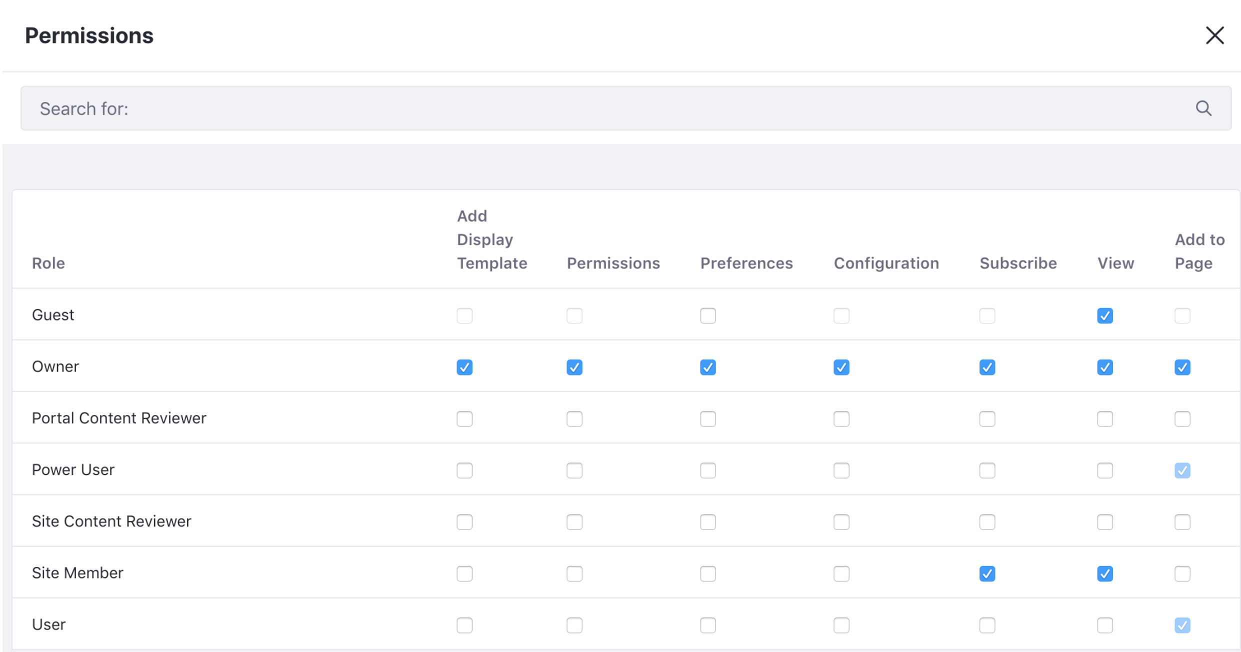

Permissions

Search for:

Role	Add Display Template	Permissions	Preferences	Configuration	Subscribe	View	Add to Page
Guest						✓	
Owner	✓	✓	✓	✓	✓	✓	✓
Portal Content Reviewer							
Power User							✓
Site Content Reviewer							
Site Member					✓	✓	
User							✓

Figure 30.12: Viewing the permissions configuration for a widget.

To set these permissions, go to the widget's *Options* icon (⋮) and click select *Permissions*. This shows you a table of Roles. Use the check boxes to grant certain permissions to different Roles. Click *Save* after you've made your selections.

Beyond this, specific permissions are generally defined for specific widgets. For example, Message Boards contains a *Ban User* permission. This makes no sense in the context of most other widgets. You'll go over permissions for specific widgets in the sections for those widgets. Next, you'll learn about widget scopes.

30.6 Widget Scope

As you learned earlier, Roles can be scoped by the instance, by Site, or by an Organization. A Role only takes effect within its scope. For example, a Message Boards Administrator Role with complete access to the Message Boards has different permissions based on the Role's scope. If it's a global Role, members have permission to administer message boards across the entire installation. If it's a Site Role, members only have permission to administer message boards within the Site where they've been assigned the Role. For Organizations with Sites, Site Roles are automatically assigned to Organization members based on the Organization Roles they have. For an Organization-scoped Message Boards administrator Role, members only have permission to administer message boards within the Site of the Organization that assigned the Role to them.

You've also heard the word *scope* refer to the data set of a widget. By default, when a widget is added to a page in a Site, it is *scoped* for that Site. This means its data belongs to that Site. If the widget is added to a page in a different Site, it employs a completely different data set. This enables you to place a Message Boards widget in one Site with one set of categories and threads, and place another Message Boards widget in different Site with a different set of categories and threads.

Scoping by Site means that you can only have one Message Boards widget per Site. If you add one Message Boards widget to a page in a Site and add another Message Boards widget to a different page in the same Site, the second Message Boards widget contains exactly the same data as the first. This is because, by default, the Message Boards widget is scoped by Site. Most of Liferay DXP's other widgets also default to being scoped by Site.

To avoid this limitation, many Liferay widgets can be scoped by page. The data sets of page-scoped widget serve a single page, not an entire Site. If you set the scope of a widget to *page* instead of *Site*, you can add any number of these widgets to different pages, and then they have different sets of data. Then you can have more than one message board per Site if you wish. Most widgets, however, default to the "native" configuration, and have their scopes set to the Site where they are placed.

Unless otherwise noted, all widgets support scoping by instance (global), Site (default), or page. This grants you some flexibility in how you want to set up your Liferay instance. You can configure the scope of an app with just a few simple steps.

1. Click the *Options* icon (⋮) in the app window.

2. Select *Configuration*.

3. Select the *Scope* tab.

4. Use the drop-down menu to set the scope.

Once you've defined a page scope for an widget, the Menu provides a *Default Scope* dropdown that allows you to select the page.

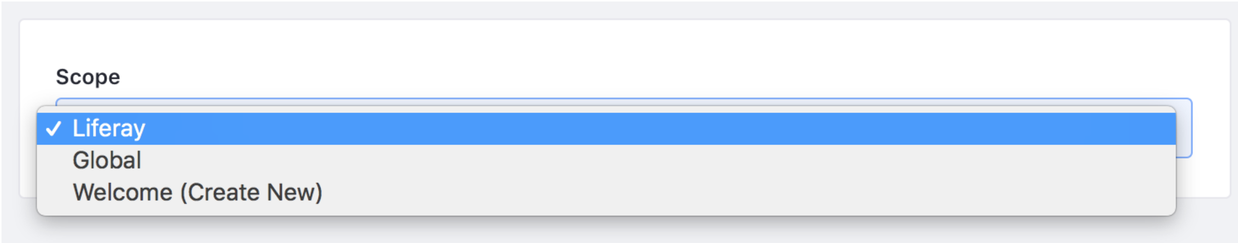

Figure 30.13: You can change the scope of your application by navigating to its Configuration menu.

That's all it takes to change the scope for a particular widget. By setting the scope to the current page, you can add as many of these widgets to a Site as you want, provided they are all added to separate pages.

Another useful feature of Liferay's apps is Configuration Templates.

30.7 Configuration Templates

Once you've configured a widget, Configuration Templates can save those settings in a reusable template. If someone goes in and changes the settings of a particular widget, it then becomes easy to revert those changes back to the original configuration template. Configuration templates are only available for widgets placed on a page. Applications available from the Product Menu do not provide configuration templates.

To create a configuration template, click the *Options* icon ( ) from the menu in the widget's title bar and select *Configuration Templates*. If widget's current settings are the ones you want to save, click the *Save Current Configuration as Template* button. If not, change the settings until it's configured the way you want it, and then click the button.

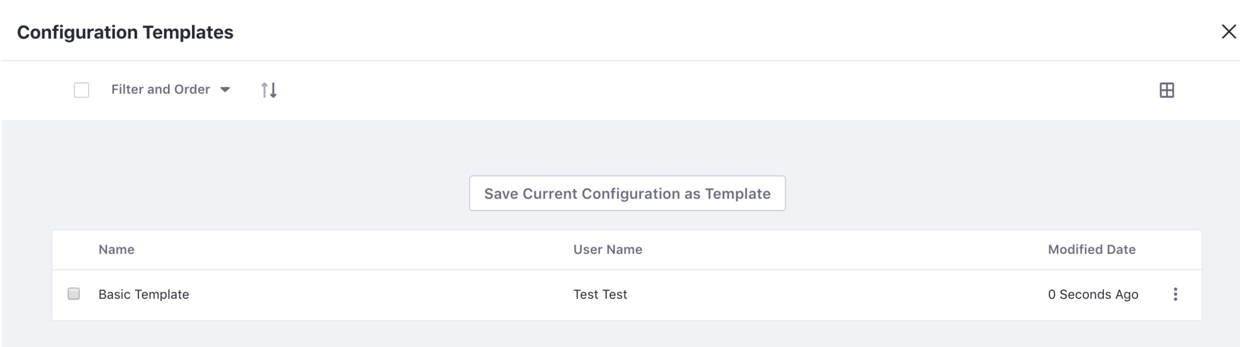

Figure 30.14: Create a configuration template to save your app's configuration settings.

There is only one field to fill out. Enter a name for your template and click *Save*. You should now see your configuration in the list. If you ever need to revert the app to these archived settings, you can click *Actions* ( ) → *Apply* next to the configuration template you want to apply.

Unless otherwise noted, all widgets in Liferay DXP support this feature. This is particularly useful for widgets that have a lot of configuration options, such as the Message Boards application.

Summary

You've now explored the configuration options available for Liferay widgets. You learned how to customize your widgets, export/import data, communicate between widgets, take advantage of different scopes, and save configuration settings. You also examined the different uses of social applications like Facebook and Netvibes for your Liferay widgets. In all, Liferay DXP gives you an abundance of options to leverage the full capability of your widgets.

Chapter 31

Staging Content for Publication

Today's enterprises generate an enormous amount of content. You can use advanced publishing tools to make the content available for users.

Staging is an important way to do this. The concept of staging is simple: you can modify your site behind the scenes and then publish all your updates in one shot. You don't want users seeing your web site change before their eyes as you're modifying it, do you? The staging environment lets you make changes to your site in a specialized *staging area* that's linked to a production environment. Typically the staging site is used only by content editors and site administrators, while the production environment is public. Content is published from staging to production all at once.

Site administrators can set up their staging environments locally or remotely. With Local Live staging, your staging environment and live environment are hosted on the same server. Remote Live staging has the staging and live environments on separate servers. You'll learn more about the differences between these two staging environments and how to enable them for your portal instance.

You can also leverage the Page Versioning feature. This feature works with both Local Live and Remote Live staging and lets site administrators create multiple variations of staged pages. This allows several different versions of sites and pages to be developed at the same time. Variations can be created, merged, and published using a Git-like versioning system. In the next section, you'll jump in to see how to enable staging.

Chapter 32

Enabling Staging

You have two different ways to set up staging: Local Live and Remote Live. Whether you enable Local Live or Remote Live staging, the interface for managing and publishing staged pages is the same.

Local Live staging lets you publish site changes quickly, since the staged and live environments are on the same server. It's also easier to switch between the staged and live environments using Local Live staging. Since the staged content, however, is stored in the same database as the production content, your server must have more resources, and the content isn't as well protected or backed up as with Remote Live staging. Also, you can't install new versions of apps for testing purposes in a Local Live staging environment, since only one version of an app can be installed at any given time on a single Liferay server.

With Remote Live staging, your staging and live environments are hosted on separate servers, so your data is separated. This lets you deploy new versions of apps and content to your staging environment without interfering with your live environment. However, publishing is slower with Remote Live staging since data must be transferred over a network. Of course, you also need more hardware to run a separate staging server.

Visit the staging environment article (Local or Remote) that most closely aligns with your goal for staging content.

32.1 Enabling Local Live Staging

Local Live staging places both your staging environment and your live environment on the same server. When it's enabled, a clone of the site is created containing copies of all of the site's existing pages. This means the staging and live environments share the same JVM, database, portlet data (depending on which portlets are selected when staging is enabled), and configurations, such as the properties set in the `portal-ext.properties` file. The cloned site becomes the staging environment and the original site becomes the live environment.

You can enable local staging for a site by navigating to the *Publishing* → Staging* menu. To get some hands-on experience with enabling Local Live staging, you can complete a brief example which creates a Local Live staging environment for your site.

1. Navigate to the Product Menu (left side) and select *Publishing* → *Staging*.

2. Select *Local Live*. You can also enable page versioning and select staged content. For more information on these options, see the Enabling Page Versioning and Staged Content article.

3. Click *Save*.

You've officially begun the staging process!

Because Local Live staging creates a clone of your site, you should only activate staging on new, clean sites. Having a few pages and some apps (like those of the example site you created) is no big deal. If you've already created a large amount of content, however, enabling staging can take a lot of time since it's a resource intensive operation. Also, if you intend to use page versioning to track the history of updates to your site, you should enable it as early as possible, *before* your site has many pages and lots of content. Your site's update history isn't saved until you enable page versioning. Page versioning requires staging (either Local Live or Remote Live) to be enabled.

If you ever need to turn off the staging environment, return back to *Staging* from the Publishing dropdown. For more information on this, see the Disabling Staging article.

Great! Now you're ready to use Local Live Staging.

32.2 Enabling Remote Live Staging

In Remote Live staging, a connection is established between the current site and another site on a remote Liferay server. The remote site becomes the live environment and the current site becomes the staging environment—an instance of Liferay used solely for staging. The remote (live) Liferay server and the local (staging) Liferay server should be completely separate systems. They should not, for example, share the same database. When Remote Live staging is enabled, all the necessary information is transferred over the network connecting the two servers. Content creators use the staging server to make their changes while the live server handles the incoming user traffic. When changes to the site are ready to be published, they are pushed over the network to the remote live server.

Before enabling Remove Live staging, ensure you've configured your Liferay server and remote server appropriately. Follow the Configuring Servers for Remote Live Staging article to do this.

You can enable remote staging for a site by navigating to the *Publishing → Staging* menu. Step through the instructions below to create a Remote Live staging environment for your site.

1. Navigate to the Product Menu (left side) and select *Publishing → Staging*.

2. Select *Remote Live*. Additional fields appear for Remote Live Connection Settings.

3. Enter your remote Liferay server's IP address into the Remote Host/IP field. This field should match the host you specified as your `tunnel.servlet.hosts.allowed` property in the `portal-ext.properties` file. If you're configuring an IPv6 address, it must contain brackets when entered into the *Remote Host/IP* field (e.g., *[0:0:0:0:0:0:0:1]*).

   If the remote Liferay server is a cluster, you can set the Remote Host/IP to the load balanced IP address of the cluster in order to increase the availability of the publishing process. See the Configuring Remote Staging in a Clustered Environment for details.

4. Enter the port on which the remote Liferay instance is running into the Remote Port field. You only need to enter a Remote Path Context if a non-root portal servlet context is being used on the remote Liferay server.

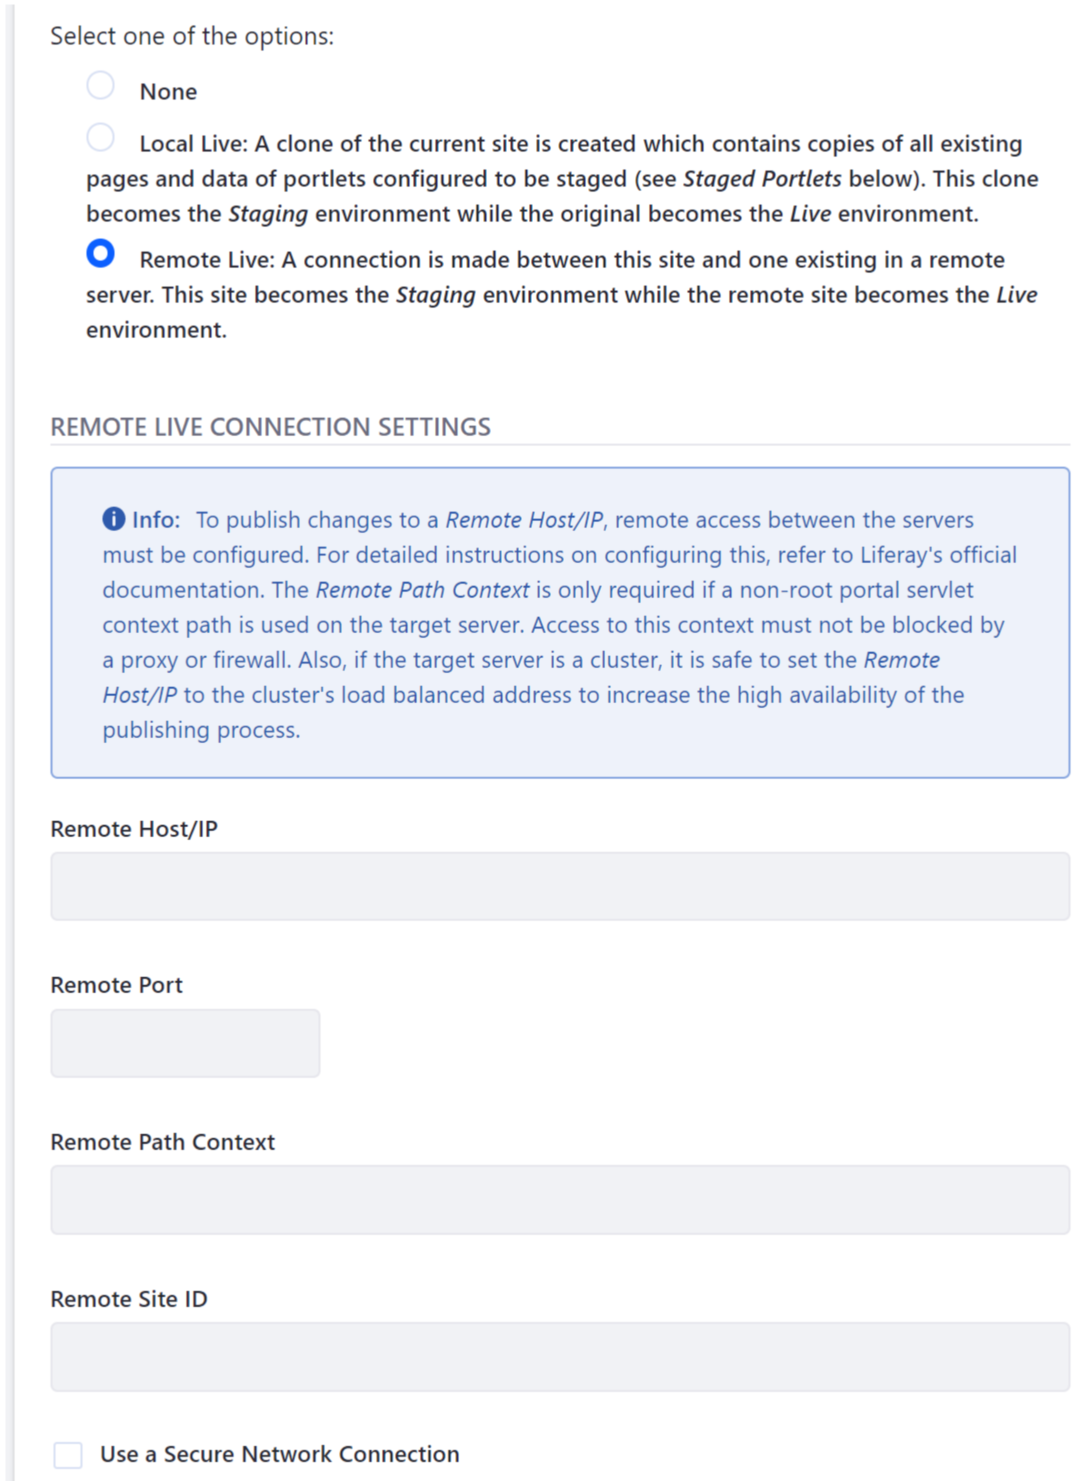

Figure 32.1: After your remote Liferay server and local Liferay server have been configured to communicate with each other, you have to specify a few Remote Live connection settings.

5. Enter the ID of the site on the remote Liferay server that's for the Live environment. If a site hasn't already been prepared on the remote Liferay server, you can log in to the remote Liferay server and create a new blank site.

   After the site has been created, note the site ID so you can enter it into the Remote Site ID field on your local Liferay server. You can find any site's ID by selecting the site's name on the Sites page of the Control Panel.

6. Check the *Use a Secure Network Connection* field to use HTTPS for the publication of pages from your local (staging) Liferay server to your remote (live) Liferay server.

7. Decide whether to enable page versioning and select staged content. For more information on these options, see the Enabling Page Versioning and Staged Content article.

8. Click *Save*.

You've officially enabled Remote Live staging!

If you fail to configure your current and remote server properly, you won't be able to enable staging and an error message appears. If you have issues, verify you've configured your servers properly.

When a user publishes changes from the local (staging) server to the remote (live) server, Liferay DXP passes the user's email address, screen name, or user ID to the remote server to perform a permission check. For a publishing operation to succeed, the operation must be performed by a user that has identical credentials and permissions on both the local (staging) and the remote (live) server. This is true regardless of whether the user attempts to publish the changes immediately or attempts to schedule the publication for later.

If only a few users should have permission to publish changes from staging to production, it's easy enough to create a few user accounts on the remote server that match a selected few on the local server. The more user accounts that you have to create, however, the more tedious this job becomes and the more likely you are to make a mistake. And you not only have to create identical user accounts, you also have to ensure that these users have identical permissions. For this reason, it's recommended that you use LDAP to copy selected user accounts from your local (staging) Liferay server to your remote (live) Liferay server. Liferay's Virtual LDAP Server application, available on Liferay Marketplace, makes this easy.

If you ever need to turn off the staging environment, return back to *Staging* from the Publishing dropdown. For more information on this, see the Disabling Staging article.

Great! Now you're ready to use Remote Live Staging.

32.3 Configuring Servers for Remote Live Staging

Before you can enable Remote Live staging for a site, you must satisfy some necessary requirements:

- Add the remote Liferay server to the current Liferay server's list of allowed servers, and vice versa.
- Specify an authentication key to be shared by your current and remote server.
- Enable each Liferay server's tunneling servlet authentication verifier.
- Update the Tunnel Auth Verifier Configuration of your remote Liferay instance.

Follow the steps below to configure your servers for Remote Live staging.

1. Add the following lines to your current Liferay server and remote Liferay server's `portal-ext.properties` file:

```
tunneling.servlet.shared.secret=[secret]
tunneling.servlet.shared.secret.hex=true
```

   Liferay DXP's use of a pre-shared key between your staging and production environments helps secure the remote publication process. It also removes the need to send the publishing user's password to the remote server for web service authentication. Using a pre-shared key creates an authorization context (permission checker) from the provided email address, screen name, or user ID *without* the user's password.

2. Specify the values for the servers' `tunneling.servlet.shared.secret` property.

   The values for these properties depend on the chosen configured encryption algorithm, since different encryption algorithms support keys of different lengths. See the HTTP Tunneling properties documentation for more information. Note that the following key lengths are supported by the available encryption algorithms:

   - AES: 128, 192, and 256 bit keys
   - Blowfish: 32 - 448 bit keys
   - DESede (Triple DES): 56, 112, or 168 bit keys (However, Liferay places an artificial limit on the minimum key length and does not support the 56 bit key length)

   To prevent potential character encoding issues, you can use one of the following two strategies:

   2a. Use hexadecimal encoding (recommended). For example, if your password was *abcdefghijklmnop*, you'd use the following settings in your `portal-ext.properties` file:

```
tunneling.servlet.shared.secret=6162636465666768696a6b6c6d6e6f70
tunneling.servlet.shared.secret.hex=true
```

   2b. Use printable ASCII characters (less secure). This degrades the password entropy.

   If you don't use hexadecimal encoding (i.e., if you use the default setting `tunneling.servlet.shared.secret.hex=` the `tunneling.servlet.shared.secret` property's value *must* be ASCII compliant.

   Once you've chosen a key, make sure the value of your current server matches the value of your remote server.

   Important: Do not share the key with any user. It is used exclusively for communication between staging and production environments. Any user with possession of the key can manage the production server, execute server-side Java code, etc.

3. Add the following line to your remote Liferay server's `portal-ext.properties` file:

```
tunnel.servlet.hosts.allowed=127.0.0.1,SERVER_IP,[STAGING_IP]
```

   The `[STAGING_IP]` value must be replaced by the staging server's IP addresses. If the server has multiple interfaces, each IP address must also be added, which would show as a source address for the http(s) requests coming from the staging server. The `SERVER_IP` constant can remain set for this property; it's automatically replaced by the Liferay server's IP addresses.

   If you're validating IPv6 addresses, you must configure the app server's JVM to not force the usage of IPv4 addresses. For example, if you're using Tomcat, add the `-Djava.net.preferIPv4Stack=false` attribute in the `$TOMCAT_HOME\bin\setenv.[bat|sh]` file.

4. Update the *TunnelAuthVerfierConfiguration* of your remote Liferay instance. To do this, navigate to the Control Panel → *Configuration* → *System Settings* → *API Authentication* → *Tunnel Authentication Verifiers*. Click */api/liferay/do* and insert the additional IP addresses you're using in the *Hosts allowed* field. Then select *Update*.

   Alternatively, you can also write this configuration into an OSGi file (e.g., `osgi/configs/com.liferay.portal.secu` `default.config`) in your Liferay DXP instance:

   ```
   enabled=true
   hostsAllowed=127.0.0.1,SERVER_IP,[Local server IP address]
   serviceAccessPolicyName=SYSTEM_USER_PASSWORD
   urlsIncludes=/api/liferay/do
   ```

5. Restart both Liferay servers after making these configuration updates. After restarting, log back in to your local Liferay instance as a site administrator.

That's all you need to do to configure Remote Live Staging! You can now enable it!
For additional information on configuring Remote Live staging, see the topics below.

Applying Patches When Using Remote Staging

When applying patches to a remote staging environment, you must apply them to all your servers. Having servers on different patch levels is not a good practice and can lead to import failures and data corruption. It is essential that all servers are updated to the same patch level to ensure remote staging works correctly.

Configuring Remote Staging's Buffer Size

Similar to Local Live staging, it is a good idea to turn remote staging on at the beginning of your site's development for good performance. When you're using Remote Live staging, and you are publishing a large amount of content, your publication could be slow and cause a large amount of network traffic. Liferay DXP's system is very fast for the amount of data being transferred over the network. This is because the data transfer is completed piecemeal, instead of one large data dump. You can control the size of data transactions by setting the following portal property in your `portal-ext.properties` file:

```
staging.remote.transfer.buffer.size
```

This property sets the file block sizes for remote staging. If a LAR file used for remote staging exceeds this size, the file will be split into multiple files prior to transmission and then reassembled on the remote server. The default buffer size is 10 megabytes.

32.4 Enabling Page Versioning and Staged Content

Enabling page versioning for a site lets site administrators work in parallel on multiple versions of the site's pages. Page versioning also maintains a history of all updates to the site from the time page versioning was enabled. Site administrators can revert to a previous version of the site at any time. This flexibility is very important in cases where a mistake is found and it's important to publish a fix quickly.

You can enable page versioning for public pages or private pages on the Staging Configuration page below the menu for selecting your staging environment (Local or Remote). If you've already enabled staging, you can navigate to the Product Menu → *Publishing* → *Staging* and click the (⋮) button and select *Staging Configuration*.

You can also choose content for the staging environment to manage on the Staging Configuration page.

Choosing content to be staged may sound self-explanatory, but content must have specific attributes in Liferay DXP to use it in a staged environment. Content or an entity should be site-scoped, so they are always part of a site; otherwise, they are not eligible for staging. For example, page-scoped entities are only eligible for staging on published pages. When scoped data is on a page (e.g., Web Content Display widget) and the page is published, the scoped data is published with it.

Liferay DXP by default supports the following content groups for staging:

- Application Display Templates
- Blogs
- Bookmarks
- Calendar
- Documents and Media
- Dynamic Data Lists
- Forms
- Knowledge Base
- Message Boards
- Mobile Device Families
- Polls
- Web Content
- Wiki

Before you activate staging, choose which of these applications' data you'd like to copy to staging. You'll learn about many of the collaboration apps listed under the Staged Portlets heading when you read the Collaboration Suite's section of articles. For now, be aware that you can enable or disable staging for any of these applications. Why might you want to enable staging for some application types but not others? In the case of collaborative apps, you probably *don't* want to enable staging since such applications are designed for user interaction. If their content were staged, you'd have to publish your site manually whenever somebody posted a message on the message boards to make that message appear on the live site. Generally, you want web content to be staged because end users aren't creating that kind of content—web content is the stuff you publish to your site. But applications like the Message Boards or Wiki should *not* be staged. Notice which applications are marked for staging by default: if you enable staging and accept the defaults, staging is *not* enabled for the collaborative apps.

The listed applications, or content groups, contain one or more specific entity. For example, selecting the Web Content application does not mean you're only selecting web content itself, but also web content folders.

Certain content types can be linked together and can reference each other on different levels. One of the responsibilities of staging is to discover and maintain these references when publishing. Site administrators and content creators have control over the process on different levels: staging can be enabled for a content group and a content group can be selected for publication.

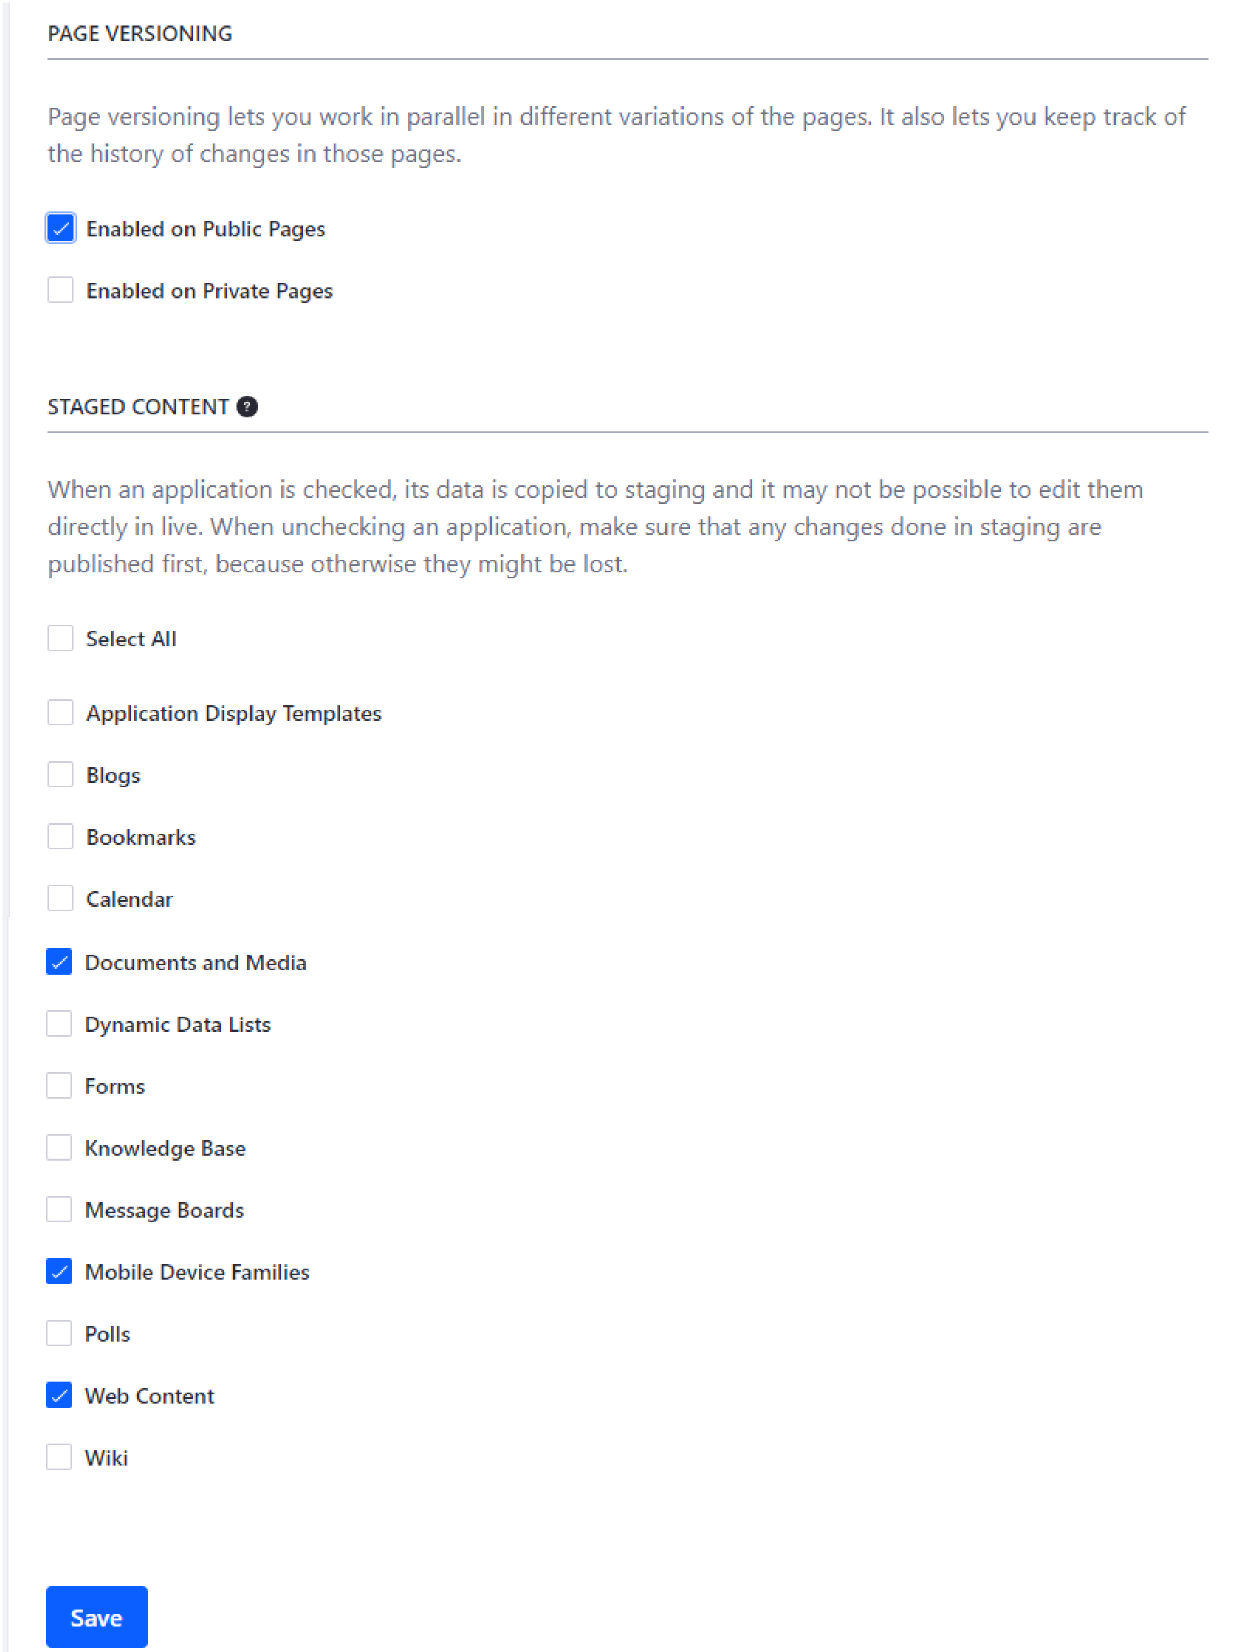

Figure 32.2: You can decide to use versioning and choose what content should be staged.

Disabled staged content types can cause unintended problems if you're referring to them on a staged site. For example, the Asset Publisher portlet and its preferences are always staged. If the content types it's set to display are not enabled for staging, the Asset Publisher can't access them on a staged site. Make sure to plan for the content types you'll need in your staged site.

Turning Staging on and off for individual portlet data could cause data inconsistencies between the staging and live sites. Because of this, it's not possible to modify the individual portlet configuration once you enable staging. In case you need adjustments later on, you must turn staging off and re-enable it with your new configuration.

Besides managing the app-specific content, Liferay DXP also has several special content types such as pages or users. For instance, pages are a part of the site and can reference other content types, but in a special way. The page references apps, which means publishing a page also implies publishing its apps. The content gives the backbone of the site; however, content alone is useless. To display content to the end user, you'll need apps as the building blocks for your site.

Before you begin exploring the Staging UI, it's important to understand the publishing process for staging to make informed decisions so you use the staging environment efficiently and effectively.

Chapter 33

Publishing Staged Content Efficiently

Now that you understand how staging works, you'll dive deeper into the publication process and some prerequisites you should follow before publishing. By understanding how the process works, you can make smart and informed decisions about how you want to publish your site's content.

33.1 Understanding the Publication Process

In simple terms, a publication is the process where all content, referenced content, apps and their preferences, pages, etc. are transferred from the staging scope to the live site. If you've enabled remote staging, this process involves network communication with another remote site. From a low level perspective, staging is an equivalence relation where entities are mirrored to a different location. From a high level perspective, the staging process happens in three phases:

1. **Export:** processes the publication configuration, which defines the site's content and apps. This phase also gathers the obligatory referenced entities that are required on the live site. Then everything according to the publication parameters is processed into the instance's own file format, and that file is stored locally or transferred to the remote live Liferay instance.

2. **Validation:** determines if it's possible to start the import process. This phase verifies the file's version and its integrity, checks for additional system information like language settings, and validates there is no missing content referenced.

3. **Import:** makes any necessary updates or additions to the site's content, layouts, and apps according to the publishing parameters. If everything is verified and correct, the staged content is published to your live site.

These phases are executed sequentially.

A crucial factor for successfully publishing staged content is data integrity. If anything is not verified during the publication process, the transactional database reverts the site back to its original state, discarding the current publication. This is a necessary action to safeguard against publishing incomplete information, which could break an otherwise well-designed live site.

If the file system is not *database-stored* (e.g., DBStore), it's not transactional and isn't reverted if a staging failure occurs. This could potentially cause a discrepancy between a file and its reference in the database. Because of this, administrators should take great care with staging the document library, making sure that regular backups of both database and file system are being maintained.

Next, you'll learn about staging best practices and prerequisites to follow for a seamless staging experience.

33.2 Planning Ahead for Staging

Staging is a complex subsystem that's flexible and scalable. Before advanced users and administrators begin using it for their site, it's important to plan ahead and remember a few tips for a seamless process. There are several factors to evaluate.

- **Content (amount, type, and structure):** Depending on the content in your site, you can turn on staging for only the necessary content types, leaving others turned off to avoid unnecessary work. Publication can also be configured to publish only certain types of content. See the Managing Content Types article for more information.

- **Hardware Environment:** You should plan your environment according to your content types. If your site operates on large images and video files, decide if a shared network drive is the best option. Storing many large images in the Document Library usually requires a faster network or local storage. If you're dealing with web content, however, these are usually smaller and take up very little disk space.

- **Customizations and Custom Logic for Your Staging Environment:** Your organization's business logic is most likely implemented in an app, and if you want to support staging for that app, you must write some code to accomplish this. You can also consider changing default UI settings by writing new JSP code if you want your staging environment's look and feel to change.

Once you've finished planning for your site, you should turn on staging at the very beginning of the site creation process. This allows the site creator to avoid waiting for huge publications that can take long periods to execute. Taking smaller steps throughout the publication process forms an iterative creative process as the site is built from the ground up, where content creators can publish their changes immediately, avoiding long wait times.

Here are some JVM/network configuration recommendations for Staging:

- 4 GB of memory

- 20 MB/s transfer rate minimum (disk)

Now that you know how the staging environment works and how to enable it for your site, you'll begin using it in the next section.

CHAPTER 34

USING THE STAGING ENVIRONMENT

After enabling staging (either Local Live or Remote Live) for a site, you'll notice additional options provided on the top Control Menu (Staging Bar) and also in the menu to the left. These new menus help you manage staged pages. Most of your page management options have also been removed; now you can't directly edit live pages. You now must use the staging environment to make changes.

!PVideo Thumbnail

Click the *Staging* button to view the staged area. Management options are restored and you can access some new options related to staging.

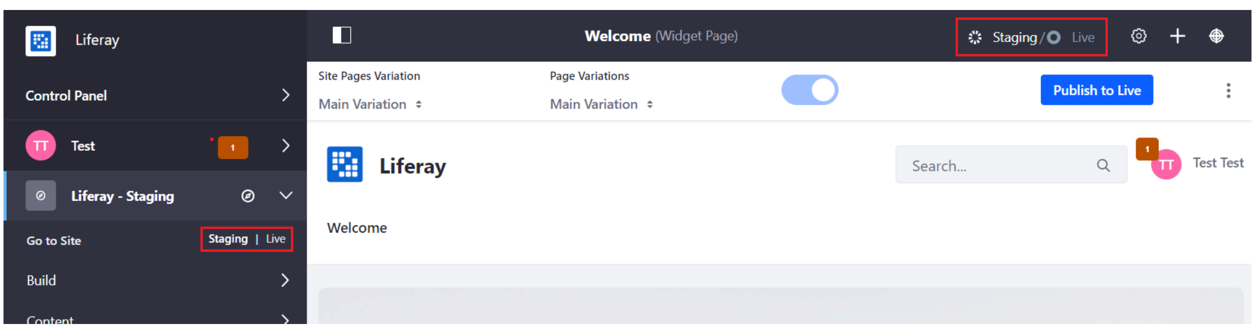

Figure 34.1: You can see the new staging options added to the top and left of your screen.

To test out the staging environment, add the Bookmarks widget and then click on *Live* from the top menu. Notice that the Bookmarks widget isn't there. That's because you've staged a change to the page but haven't published that change yet to the live site.

Next, you'll learn the basics of staging content.

!VVideo Tutorial

34.1 Staging Content

When you're on the staged site, you have several options in the Staging Bar to help start your staging conquest.

Site Pages Variation: lets you work in parallel on multiple versions of a staged site page. You can choose the site page variation from the dropdown menu or manage them from the *Options* icon ( ) in the Staging Bar. See the Using Site Pages Variations article for more information.

Page Variations: lets you work in parallel on multiple versions of a staged page. You can choose the page variation from the dropdown menu or manage them from the *Options* icon ( ) in the Staging Bar. See the Using Site Pages Variations article for more information.

Undo/Redo: steps back/forward through recent changes to a page, which can save you the time of manually adding or removing apps if you make a mistake. To access *Undo/Redo*, select the *Options* icon ( ) in the Staging Bar.

History: shows the list of revisions of the page, based on publication dates. You can go to any change in the revision history and see how the pages looked at that point. To access *History*, select the *Options* icon ( ) in the Staging Bar.

Ready for Publication: After you're done making changes to the staged page, click this button. The status of the page changes from *Draft* to *Ready for Publication* and any changes you've made can be published to the Live Site. When you publish a page to live, only the version *Marked as Ready for Publication* is published.

Now you'll step through a brief example for using the Staging Bar to stage and publish content.

1. On the staged site, navigate to the Product Menu and select *Content* → *Web Content*.

2. Create a Basic Web Content article and save it.

3. Go back to your staged site's main page and navigate to the *Add* ( ) → *Widgets* → *Content Management* menu and drag the *Web Content Display* widget to the page.

4. Select the web content article you created to display in the new widget.

5. Select the *Ready for Publication* button to confirm you're ready to publish the content from the staged site to the live site. This prepares the staged content for publication. If workflow is enabled for any new resource, the resource must go through the workflow process before it can be published to the live site.

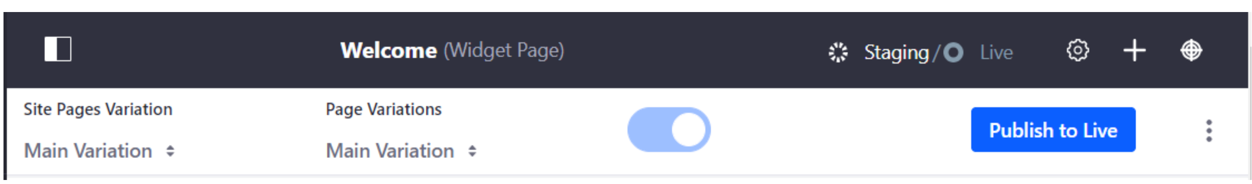

Figure 34.2: The staging toolbar indicates whether you're able to publish to the live site.

6. Click the *Publish to Live* button. A pop-up window appears with configuration options for your publication.

7. Enter the name of your publication.

8. Observe the changes listed in the menu. This lists the changed content planned for publication.

9. Click the *Publish to Live* button to publish your staged results to the live site.

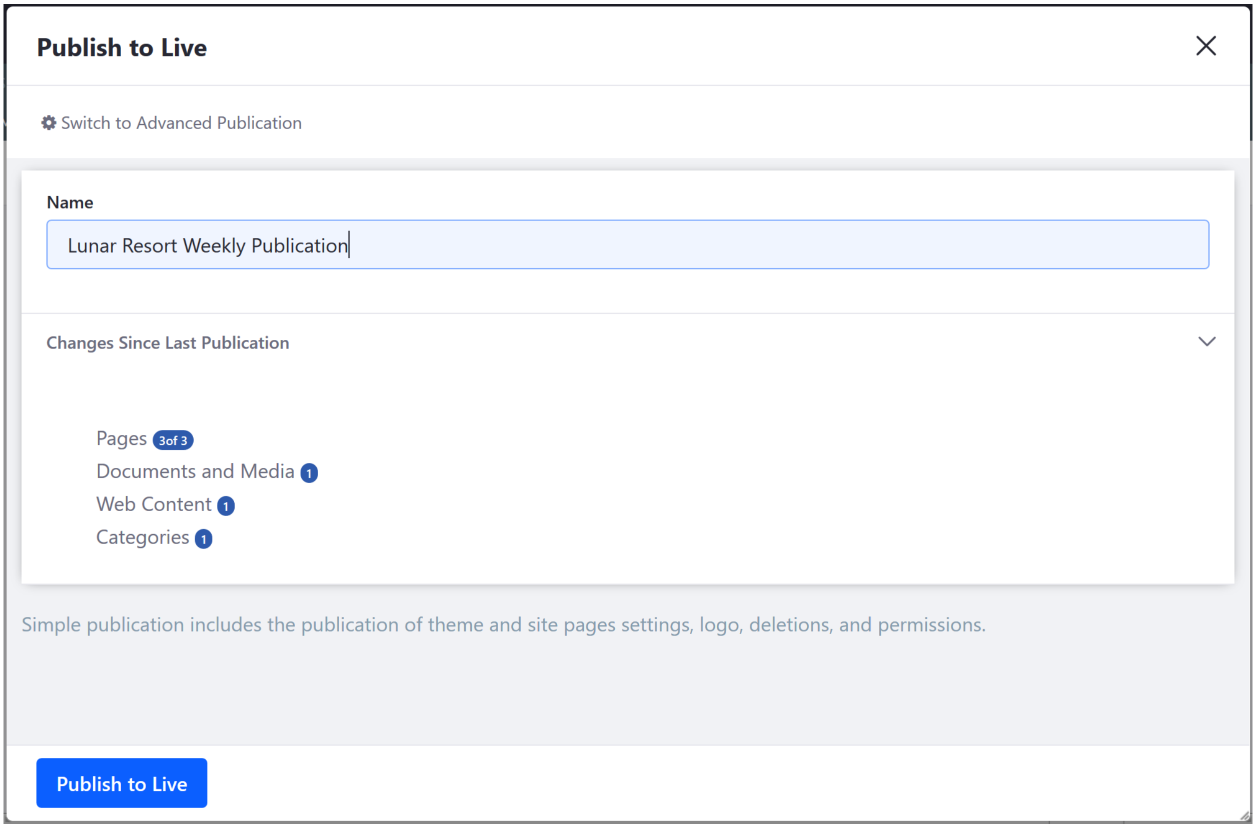

Figure 34.3: The Simple Publication menu displays the changes since last publication and a way to name your publication.

Awesome! You've created content on the staged site and published it to your live site. It's now available for all your site users to see!

Note: Although publishing content is the more well-known function, publishing a portlet is also a viable option. You can publish portlets residing in the Control Panel and on pages. For example, you can modify a portlet's title and publish the change to live. This is possible because portlet configurations are always staged.

To publish a portlet that is on a page, you must publish the page first.

This example explored the Simple Publication menu. If your publication requires more advanced configuration like specifying specific content, dates, pages, etc., you should click the *Switch to Advanced Publication* button. You'll explore the more advanced configuration options next.

34.2 Advanced Publication with Staging

Once you've finished your changes on the staged site and want to publish them, select the *Publish to Live* button from the Staging Bar. To configure advanced publication options, select the *Switch to Advanced Publication* button. Opening the Advanced Publication menu presents options for scheduling a time to publish your content, editing the pages/content to include in the publication, managing permissions, etc. This lets you perform advanced editing to your publication process.

Date

You have two options for the Date category:

Now: immediately pushes any changes to the live site.

Schedule: set a specific date to publish or to set up recurring publishing. You could use this to publish all changes made during the week every Monday morning without any further intervention.

These options let you plan staging schedules in advance.

Pages

This area of the menu gives you the option to choose which pages to include when you publish. You can choose the page group (Public or Private) to publish by selecting the *Change to Public Pages* or *Change to Private Pages*. You cannot publish both at the same time; you must complete their publication processes separately if you want to publish both page groups.

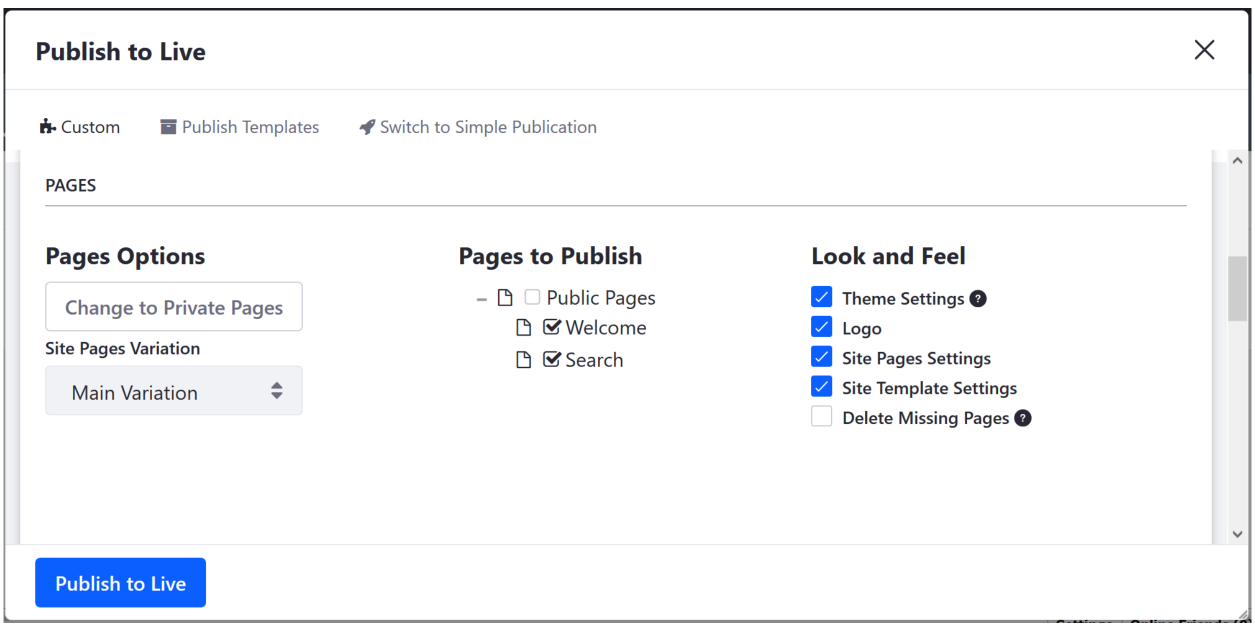

Figure 34.4: You have several ways to specify the pages you want included in your publication.

Note: If you're publishing pages with a custom theme, you must check the *Theme Settings* option under the Look and Feel heading for your staging configuration. If it's not checked, the default theme is always applied.

You can also choose specific pages to publish, and the look and feel of those pages.

Content

The Content area lets you select the content to be published. If you choose a page to be published from the Pages menu, the portlets and their references are always published, even if you specify differently in the Content section.

There are other filtering sub-options for certain content types. You first must choose what content to publish based on date. Specifying a date range lets you choose content to publish based on when it was created or last modified. Select the option that best fits your workflow. The available options are described in more detail below:

All: publishes all content regardless of its creation or last modification date.

From Last Publish Date: publishes content that was created or modified since the last publish date. This is the default option.

Date Range: publishes content based on a specified date range. You can set a start and end date/time window. The content created or modified within that window of time is published.

Last: publishes content based on a set amount of time since the current time. For example, you can set the date range to the past 48 hours, starting from the current time.

Under the date options are the different types of content that can be published. This list is populated based on the provided date range. For example, if at least one article is created or modified in the given date range, a Web Content section appears in the list, and the number of articles is shown next to the Web Content label. Otherwise, the Web Content section is absent.

The *Categories* content type is not dependent on the date range and is always shown in the list.

Note: Since some content types are meant for the end user and aren't supported in staging (e.g. comments, ratings, and custom fields), they can only be added to the live site and cannot be removed.

Unchecking the checkbox next to a certain content type excludes it from the current publication to the live site.

Some of the content types in the list, like Web Content and Documents and Media, have further filtering options. For instance, when the Web Content section is present and checked, it shows a comma-separated list of related items to be published, including the articles themselves. A sample list of related items for web content might look like this: *Web Content(12), Structures(3), Referenced Content, Include Always, Version History*. You can remove items by clicking the *Change* button next to the list.

See the Managing Content Types in Staging article for more information on managing content during the publication process.

Deletions

This portion of the menu lets you delete two things:

- portlet metadata before publishing
- operations performed for content types.

You have two options to manage for deletions:

Delete Application Data Before Importing: all data created by the application is deleted before the import process. Ensure you understand the ramifications of this option before selecting it. Some applications in other pages may reference this data. This process cannot be undone. If you are unsure, complete an export first.

Replicate Individual Deletions: operations performed for content in the staging environment are replicated to the target site.

Permissions

This area lets you include permissions for the pages and portlets when the changes are published. Select the *Publish Permissions* checkbox to enable this functionality.

Once you're finished configuring you advanced publication, select *Publish to Live* to publish or schedule your publication.

34.3 Managing Content Types in Staging

When managing content in Staging's Advanced Publication menu, there are several factors to consider when preparing your content for publication. As described in Advanced Publication with Staging, you can navigate to the Content area of the Advanced Publication menu to select content you want to publish. There are options attached to each content group (e.g., Web Content) that you can manage too.

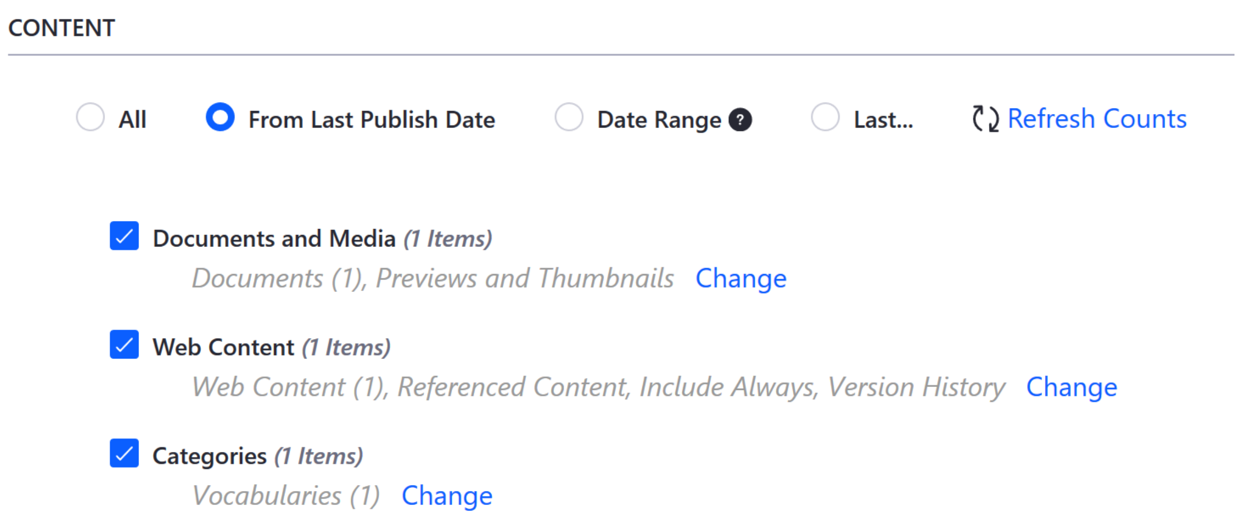

Figure 34.5: Click the *Change* button for a content group to manage its specific content.

You'll learn about some of these options and their best practices next.

Referenced Content

This is represented by

- Structures and templates included in web content.
- Documents and Media files (e.g., images) included in web content.
- etc.

You can exclude some of this content during publication or export to speed up the process. These references are validated during the publication process or an import, so the images must be published or imported first.

Version History

Web content tends to be updated frequently, often more so than other kinds of content. Sometimes this can result in high numbers of versions. If there are hundreds of versions, it can take a long time to publish these articles. You can bypass this by choosing to not publish the *Version History* (i.e., the past versions of the web content articles to be published). If you disable this, only the last **approved** version of each web content article is published to Live. This can significantly speed up the publication process.

You can set this option globally. If you navigate to the Control Panel → *Configuration* → *System Settings* → *Web Content Administration*, you can toggle the *Publish version history by default* checkbox. This sets the default behavior. When publishing content, it is selected by default, so site administrators must manually uncheck the *Version History* box to publish only the latest approved version of web content articles. To change the default behavior, enable the checkbox in System Settings.

Previews and Thumbnails

Previews and thumbnails are generated automatically for documents. Disabling this, though, can greatly increase your publishing speed in some cases. You should be careful about publishing previews and thumbnails to the live site.

Imagine a scenario where a site has approximately 4000 images or documents. If the previews and thumbnails are turned on, this could end up in 28000 physical files on the disk. If staging is set up to publish the previews and thumbnails, this would mean that instead of taking care of the 4000 images, it would process seven times more files! If you still want to use the previews on your live environment, you can set up that Liferay instance to generate them automatically.

It depends on your environment for whether you can use the publishing of the previews and thumbnails. Publishing them is a heavy operation, and you must also transfer the LAR file over the network if you use remote staging. If you decide to generate them on the live site, understand that this could take some time, since it's a CPU intense operation.

Vocabularies

When working within a site, a user may select vocabularies from both the current site as well as the global site. While this doesn't pose an issue when creating content, it can cause issues when publishing.

For environments that use both global and local vocabularies, note that global vocabularies must be published to the live site through global site staging. One way to avoid confusion with vocabularies is to keep all vocabularies local or global.

If both must be used, you can resolve the issue by ensuring that dependencies (e.g., categories and vocabularies) are published before publishing the site that depends on them (whether the dependencies are local or global).

Assets like tags, categories, structures, templates, application display templates, document types, and dynamic data lists can also be shared by a parent to its child sites. In this case, ensure that the ancestor's dependencies are published before the site in question.

Deletions

The Staging framework gathers deletions (including trashed entities) in a site. These deletions can be published to clean up the live site. If you plan to process it later, or if it's not a problem to have lingering data on live, this can be turned off as well to save execution time during the process.

34.4 Staging Processes and Templates

When you make a staging publication, it's captured as a staging process and stored for future reference. You can manage these processes by navigating to the *Staging* option located in the Product Menu's Publishing tab. From there, you'll see a list of staging processes that have been completed. You can relaunch or clear any of these publications by clicking the *Actions* button (⋮) next to a process. If you click the *Scheduled* tab from above, you'll find staging processes that you've scheduled for future publication dates.

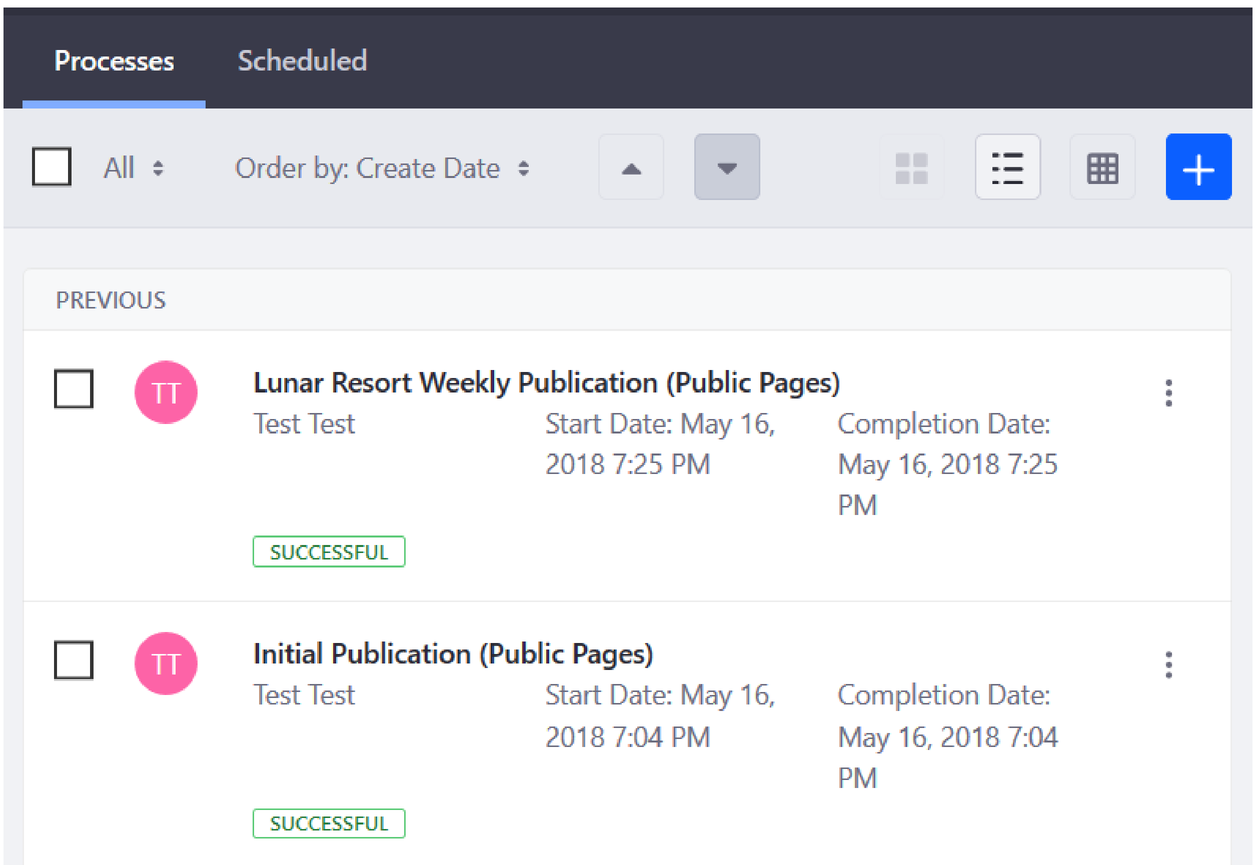

Figure 34.6: Your staging processes can be viewed at any time.

If you find yourself repeatedly creating similar staging processes, you should think about using Publish Templates.

Instead of manually having to customize a publication process every time you're looking to publish pages/content, you can use a publish template. With publish templates, you can select a custom template and immediately publish with the options you configured.

Follow the steps below to create and use a publish template.

1. Select the *Options* icon () from the top right corner of the Staging screen and select *Publish Templates*.
2. Click the *Add* button () and assign the template a name and description, and then fill out the configuration options as you would during a custom publication process.
3. Save your publish template. It's available to use from the *Publish Templates* tab in the *Publish to Live* menu's Advanced Publication area.
4. To use the template, click the *Actions* button ( ) next to the template and select *Publish to Live*.

   This automatically sets the options for publishing pages and their content. All you have to do is give the publication process a name. Once you confirm the configuration settings, your staging settings are published.

Note: When staging is enabled, the options available from the *Publishing* tab are modified. When in the Live environment, you can only access the *Export* feature. When in the Staging environment, you can only access the *Import* and *Staging* features. The disabled features for each environment don't make sense in that context. For example, you shouldn't be able to import content when in the live environment; it must be imported into the staged environment and then published before it is available in the live site.

Now you know how to reference stored/scheduled staging processes and create publish templates to streamline publication.

Chapter 35

Disabling Staging

Disabling staging doesn't take a lot of steps, but should not be done lightly. It's important to know the consequences of turning the staging environment off so you can decide if your circumstances really warrant it.

The consequences for disabling Local Live and Remote Live staging are slightly different, so you'll learn about both.

35.1 Disabling Local Live Staging

Conceptually, the live site is the final approved version of your site, whereas the staging site is a temporary workspace containing information that is not finalized.

Disabling local staging completely removes the staging environment, which means all content not published to your live site is erased. Therefore, before disabling staging, you must ensure all necessary information on the staged site is published or preserved elsewhere.

Keep in mind that draft content types are not published, so they can be lost too.

When you enabled staging there was an initial publication. Disabling staging does not start a publication; the staging site is deleted. If the staged site contains a large amount of content, however, those deletions could take a substantial amount of time to process. For this reason, don't disable staging when your portal instance is busy.

35.2 Disabling Remote Live Staging

Disabling remote staging does not delete the staged site; it only disables the connection between the live site and remote staging site. This means no data is deleted from the live or remote staging sites after the connection is disabled. Since no data is erased and no processes are started, disabling remote staging is almost instantaneous.

When Remote Live staging is enabled, certain information (e.g., which portlet is being staged) is recorded on both the live and staged sites. For this reason, when you disable remote staging, you must ensure the live site is still accessible so both sides can communicate that it's disabled. Do not shut down your live site and then attempt to disable remote staging from your staged site; this results in errors.

If there's ever a lost network connection between the remote staged site and the live site, a message appears, informing you of the error and a way to forcibly disable staging. This is only an option for the staged site; executing this option erases the staged site's staging information—not the content. On the contrary, the live site remains in a locked state. A possible workaround is to create a new live site and import content to it, if necessary.

35.3 Steps to Disable Staging

Follow the steps below to disable Local Live or Remote Live staging:

1. Navigate to the *Publishing* → *Staging* option, which is only available from the staged site.
2. Click the *Options* icon () from the upper right corner of the page and select *Staging Configuration*.
3. Select the *None* radio button and click *Save*.

That's it! Your staging environment is now turned off.

Chapter 36

Publishing Single Assets From a Staged Site

Sometimes, stepping through the entire publication process is not necessary and can be overkill. For example,

- What if you created a web content article and want only to publish it and nothing else?
- What if you want to publish a new folder of articles and their dependencies, but don't want the hassle of the publication process?
- What if you found a typo in a document and want to fix it quickly and push to the live site?

You're in luck! You can publish certain single assets from the staged site to the live site without creating a staging publication process, from their respective app menus:

- Web Content
  - Web Content
  - Folder
- Documents and Media
  - Document
  - Folder
  - Shortcut
  - Document Type
- Blogs
  - Blog
- Bookmarks
  - Bookmark
  - Folder

Important: Single asset publication is not supported for page-scoped content.

Note: When publishing a Web Content or Bookmarks folder, their respective entries and subfolders are included. Publishing a Documents and Media folder works the same way, but also includes shortcuts.

You'll step through an example to see how this is done.

1. Make sure the Staging framework is enabled and you're on the staged site.

2. Create a Web Content Article in the Product Menu's *Content* → *Web Content* menu.

3. Once you've saved the new Web Content Article, select its the *Actions* button ( ) next to the article and select *Publish to Live*.

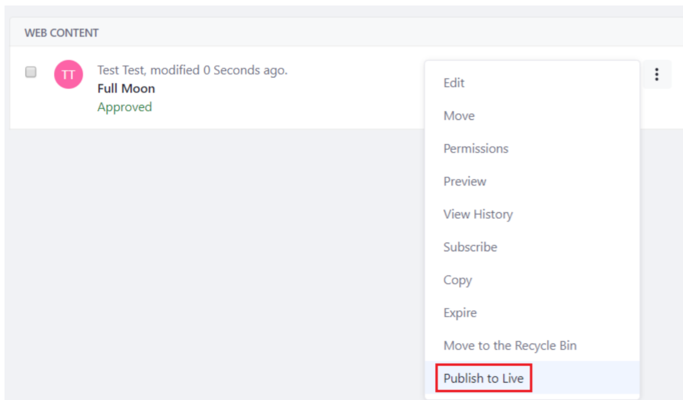

Figure 36.1: You can publish the single web content article to the live site.

4. You're presented a Process Details page where you can view the progress of your single asset publication request. Ensure the Web Content Article is published successfully.

```
**Note:** Sometimes the publication process doesn't start immediately (e.g.,
if there's another publication running). You can check a specific asset's
publication progress by navigating to the *Options*
(![Options](./images/icon-options.png)) &rarr; *Staging* &rarr;
 *Current and Previous* tab in its Site Admin app.
```

There you have it! If you navigate to your live site's Web Content section, the new article is available.

Similar to the regular staging publication process, your single asset publications also include associated dependencies. For example, if your web content article contains an image, custom structure, and custom template, they are all published together. The same concept applies for folders—if you publish a folder containing several web content articles, all the articles and their associated dependencies are published too.

By default, only those with permissions to publish widgets can publish single assets. Follow the steps below to modify these permissions for a Role:

1. Navigate to the Control Panel → *Users* → *Roles*.

2. Select the Role you're updating.

3. Click the *Define Permissions* tab.

4. In the left menu, navigate to *Control Panel* → *Sites* → *Sites*.

5. Under the Resource Permissions heading, select the *Export/Import Application Info* option.

Granting a Role the *Publish Staging* permission lets users assigned the Role publish single assets. Managing Permissions article for more information.

Great! Now you know how to publish single assets and manage the permissions for who can do it.

Chapter 37

Organizing Pages for Staging

Say you're working on a product-oriented Site where several major changes are required for a page or a set of pages over a short period of time. You must work on multiple versions of the Site at the same time to ensure everything has been properly reviewed before it goes live. With staging, you can do this using *Page Variations*.

In this section, you'll explore page variations and how they're useful.

37.1 Using Multi and Single Page Variations

There are two page variation options available from the Staging Bar:

- *Site Pages Variation:* Different variations for a set of Site pages. For instance, you could use this if you had three separate pages and wanted to modify these pages while keeping them together as a set.
- *Page Variations:* Different variations for a single page. A page variation can only exist inside a Site pages variation.

You must enable page versioning in the Staging Configuration menu before the above options are available in the Staging Bar. You'll learn more about this later.

Variations only affect pages and not the content, which means all the existing content in your staging Site is shared by all your variations. The content, however, can be displayed in many different ways for each page variation. For example, in different Site page variations, you can have different logos, a different look and feel for your pages, different applications on these pages, different configuration of these applications, and even different pages. One page can exist in just one Site page variation or in several of them. Modifying the layout type (e.g., Layout, Panel, Embedded, etc.) or friendly URL of a page, however, **does** affect every Site page variation. For example, if a page template is modified, those modifications are propagated to the pages configured to inherit changes from the template, overriding Staging's Page Variations and Site Pages Variations.

Note: Page templates are not recognized by the Staging framework. This means that existing page templates are not viewable or editable on a staged Site. If they're created on a staged Site, they cannot be preserved once Staging is disabled. You can, however, export and import page templates.

You'll learn about enabling page versioning next.

Enabling Page Versioning

Page Versioning is enabled on the Staging Configuration screen when first enabling staging.

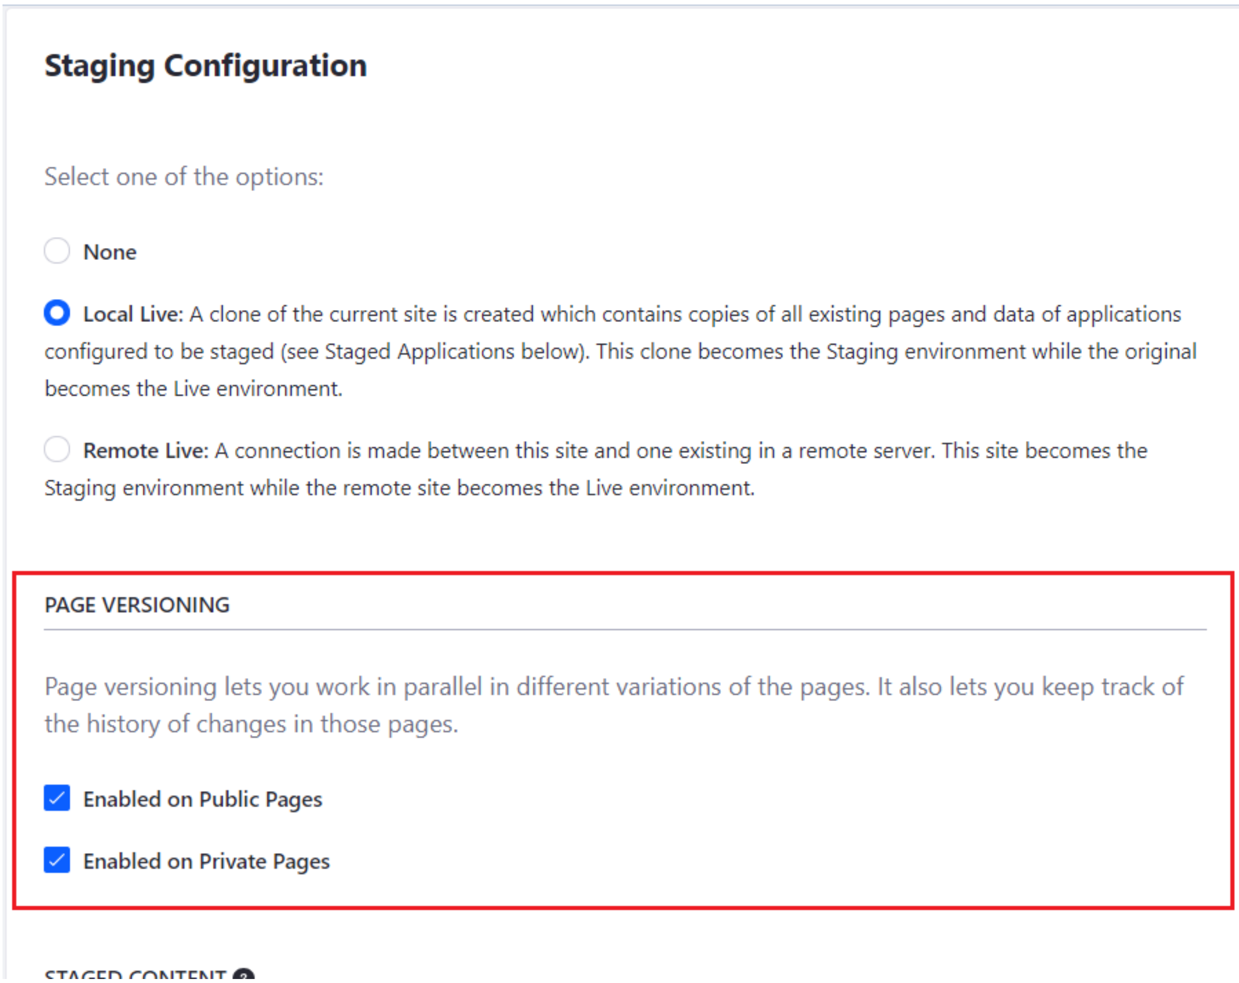

Figure 37.1: You can enable page versioning for public and/or private pages.

You can enable page versioning for public and private pages. When page versioning is enabled, the page variation options are available in the Staging Bar. By default, you only have one Site pages variation and page variation which are both called *Main Variation*.

If you did not enable page versioning during the initial setup of your staging environment, navigate to the Product Menu → *Publishing* → *Staging* → *Options* () → *Staging Configuration*. You can enable the page versioning options there.

Using Site Pages Variations

Site pages variations are useful when you must plan multiple page sets for your Site at once. For example, consider this scenario:

If there were separate teams in your company that needed to create three drastically different page sets for your Site at the same time, they would need to create three Site pages variations. For example,

- The marketing team can give your Site a completely different look and feel for the Holidays.
- The product management team can work on a version that is planned to publish on the first of the year for a new product launch.
- The developer relations team can work on a version that is planned to publish on the upcoming Hack-a-thon day.

With this use case, having a Site pages variation for each planned page set allows three ideas to be fully planned and implemented before publication.

Another option for this scenario is to let the product management team experiment with two different ideas for the home page of the Site. This can be done by creating several page variations within the current Site pages variation of their product launch Site. You'll learn more about page variations later.

Once you've created a Site pages variation, you can now navigate to its home page and change the logo, apply a new theme, move applications around, change the order of the pages, configure different apps, and more. The other variations aren't affected. You can even delete existing pages or add new ones (remember to *Mark as Ready for Publication* when you are finished with your changes). When you delete a page, it is deleted only in the current variation. The same happens when you add a new page. If you try to access a page that was deleted in the current variation, Liferay informs you this page is not *enabled* in this variation and you must enable it.

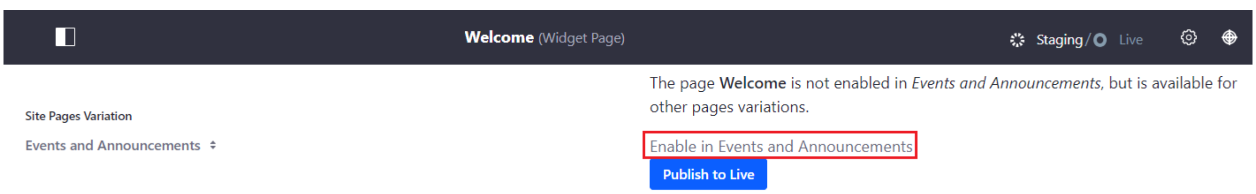

Figure 37.2: Select the *Enable* button to create a missing page in the current Site pages variation.

To publish a variation to the live Site, click on *Publish to Live* in the staging menu and then select *Publish to Live*. Publications can also be scheduled independently for different variations. For example, you could have a variation called *Mondays* which is published to the live Site every Monday and another one called *Day 1* which is published to the live Site every first day of each month.

Using Page Variations

You can also have variations for a single page inside a Site pages variation, which lets you work in parallel on different versions of a page. For example, you might work on two different proposals for the design of the home page for a Holidays Site pages variation. Page variations only exist inside a Site pages variation.

Once you've created a page variation, you can choose it from the dropdown menu on the Staging Bar. You can always switch between different variations by clicking on them from the Staging Bar.

Once you've modified the page variation to the way you want, mark it as *Ready for Publication*. Only one page variation can be marked as ready for publication and that is the one that gets published to the live Site. To publish a variation to the live Site, click on *Publish to Live* in the staging menu and then select *Publish to Live*.

Managing Variation Permissions

It's also possible to set permissions on each variation, so certain users have access to manage some, but not all variations. To do this,

1. Navigate to the Staging Bar's *Options* button () and select the variation type you want to configure.

2. Select the desired variation's *Actions* button ( ) and select *Permissions*.

3. Configure the variation's permissions and then click *Save*.

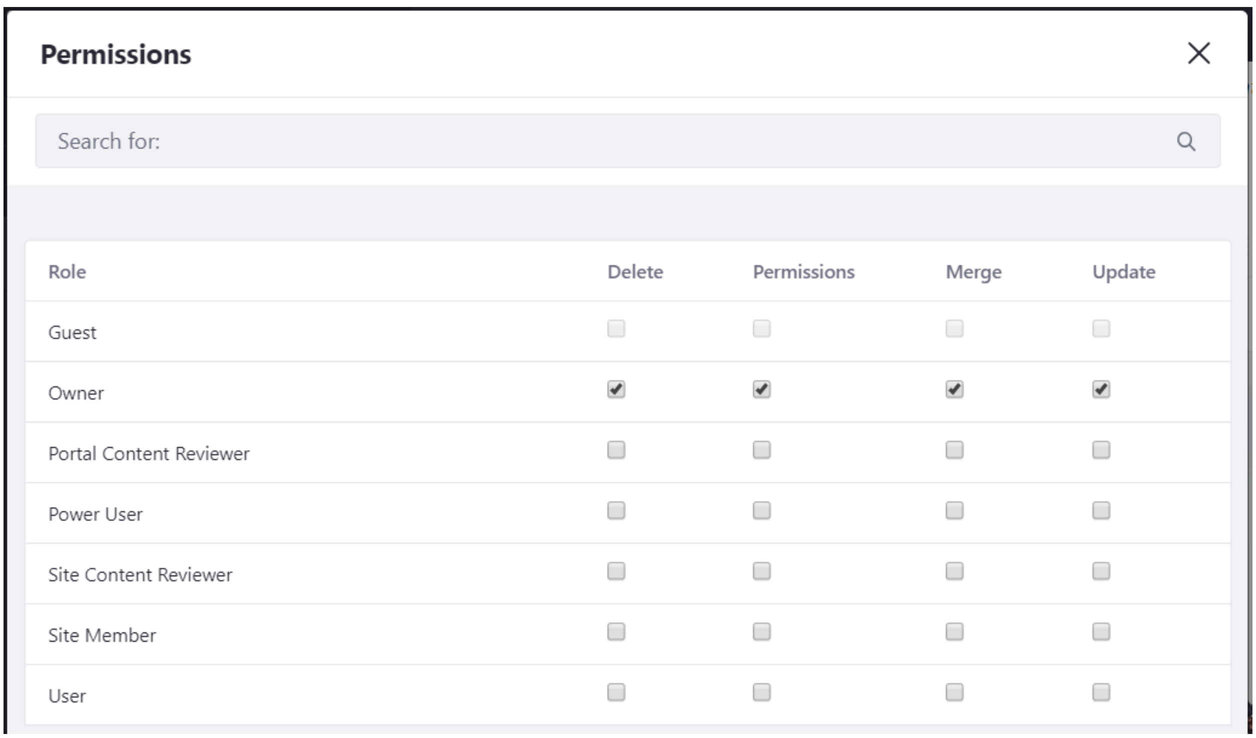

Figure 37.3: Configure the roles that can access and modify your variation.

Awesome! You now know how to manage variation permissions!

37.2 Creating Multi and Single Page Variations

You can create two types of page variations:

- Site Pages Variation (multiple pages)
- Page Variation (single page)

You can learn more about these variations in the Using Multi and Single Page Variations. As an example, you'll step through creating a Site pages variation next. Both variation processes, however, are similar.

1. Select the *Options* icon ( ) in the Staging Bar and select the variation option. For example, select the *Site Pages Variation* option. This brings you to a list of the existing Site page variations for your Site.

2. Click *Add Site Pages Variation* to create a new one.

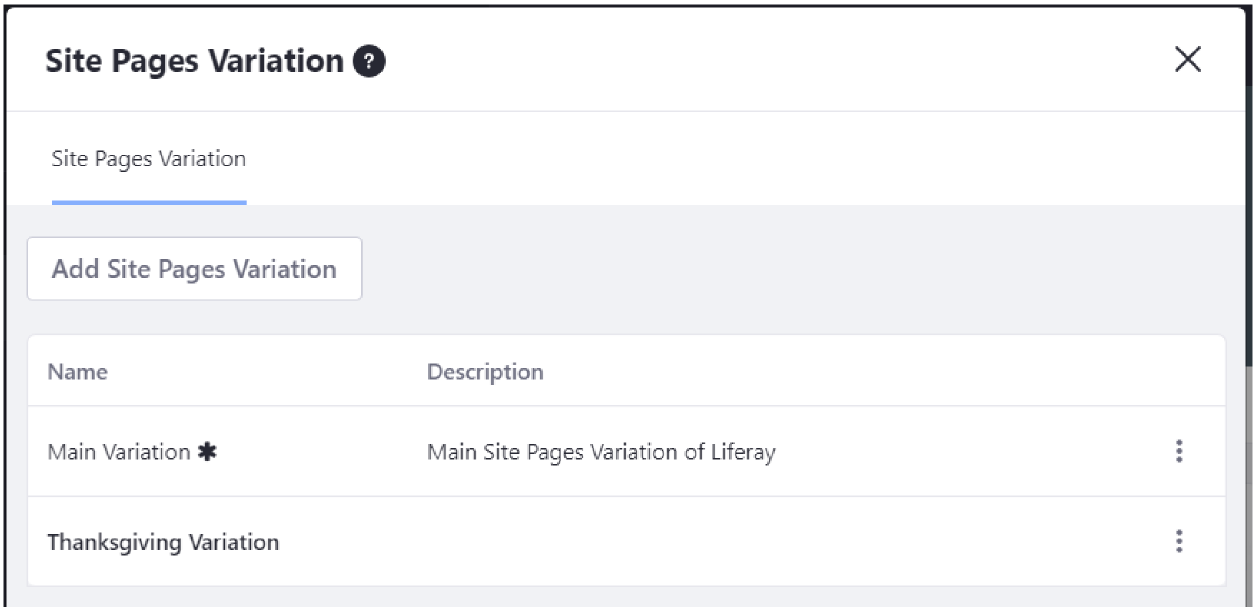

Figure 37.4: When selecting the *Site Pages Variation* link from the Staging Bar, you're able to add and manage your Site pages variations.

3. Set a name and description for your new Site pages variation.

4. Set how you want your variation created. From the *Copy Pages from Site Page Variation* field, you can copy content from an existing variation to create your new one. There are several options to choose in this selector.

   All Site Pages Variations: Creates a new variation that contains the last version marked as ready for publication from any single page existing in any other variation.

   None (Empty Site Pages Variation): Creates a new, empty variation.

   [Existing Variations]: Creates a new Site page variation that contains only the last version of all the pages that exist in a specific variation (e.g., *Main Variation*). The current variation must be marked as ready for publication.

   The copy option is not available when creating a page variation. A new page variation is a copy of the current Site pages variation.

5. Click *Add* to create the Site pages variation.

Note: You can rename any variation after it's created, if necessary. For example, edit the Main Variation and change its name to something that makes more sense in your Site, such as *Basic*, *Master*, or *Regular*.

Awesome! Your Site pages variation is created and available for modification.

37.3 Merging Site Pages Variations

Another powerful feature of Staging's Page Versioning framework is the possibility of *merging* Site Pages Variations. To merge two Site Pages Variations, follow the instructions below.

1. Click the Staging Bar's *Options* button (⋮) and select *Site Pages Variation*.
2. Click the Site Pages Variation's *Actions* button (⋮) you want to use as the base for merging and select *Merge*.
3. Select the Site Pages Variation to merge on top of the base Site Pages Variation.

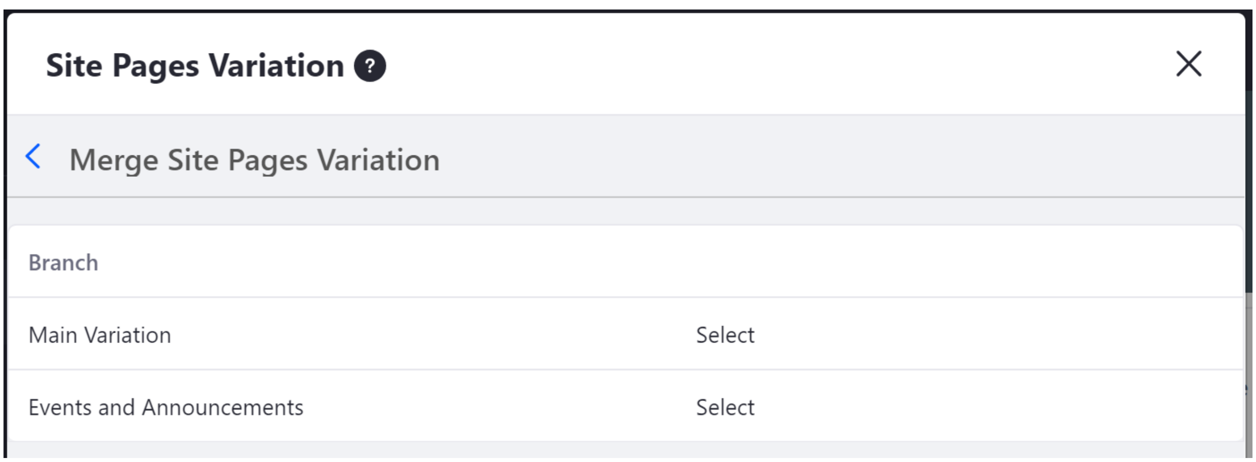

Figure 37.5: Select the site pages variation you'd like to merge with your base variation.

Merging works like this:

- New pages that don't exist in the base variation are added.
- If a page exists in both Site Pages Variations, and at least one version of the page was marked as ready for publication, the latest version marked as ready is added as a new page variation in the target page of the base variation. Note that older versions or page variations not marked as ready for publication aren't copied. Merging can be executed, however, as many times as needed and creates the needed page variations in the appropriate page of the base site pages variation.
- Merging does not affect content and doesn't overwrite anything in the base variation; it adds more versions, pages, and page variations as needed.

Great! You've merged site pages variations!

Chapter 38

Managing Permissions

The staging environment has many different options for building and managing a Site and its pages. Sometimes administrators want to limit access to staging's subset of powerful features. A Role with permissions can accomplish this. To create/modify a Role, complete the following steps:

1. Navigate to the *Control Panel* → *Users* → *Roles*.
2. To create a new Role, select the *Add* button () and complete the New Role menu. Once you have a new Role created, or you've decided on the Role you want to modify, select the Role's *Actions* icon ( ) and select *Edit*.
3. From the top menu, select *Define Permissions*.

The most obvious permissions for staging are the general permissions that look similar to the permissions for most Liferay apps. These permissions are in the *Site Administration* → *Publishing* → *Staging* section of the Define Permissions menu. They include

- *Access in Site Administration*
- *Add to Page*
- *Configuration*
- *Permissions*
- *Preferences*
- *View*

Also, there are some Site resource permissions that deal directly with staging. These permissions are in the *Control Panel* → *Sites* → *Sites* section in the Define Permissions menu. The relevant Site resource permissions related to staging are listed below:

Add Page Variation: Hides/shows the *Add Page Variation* button on the Staging Bar's Manage Page Variations screen.

Add Site Pages Variation: Hides/shows the *Add Site Pages Variation* button on the Staging Bar's Manage Site Page Variations screen.

Export/Import Application Info: If the Publish Staging permission is not granted, hides/shows the application level Export/Import menu. The Configuration permission for the Export/Import app is also required.

Export/Import Pages: If the Publish Staging permission is not granted, hides/shows the Export/Import app in the Site Administration menu.

Manage Staging: Hides/shows the Staging Configuration menu in the Site Administration → *Publishing* → *Staging* → *Options* (⋮) menu.

Publish Application Info: Hides/shows the application level Staging menu.

Publish Staging: Hides/shows the *Publish to Live* button on the Staging Bar and hides/shows the *Add Staging Process* button (+) in the Site Administration menu's Staging app. This permission automatically applies the *Export/Import Application Info*, *Export/Import Pages*, and *Publish Application Info* permission functionality regardless of whether they're unselected.

View Staging: If Publish Staging, Manage Pages, Manage Staging, or Update permissions are not granted, hides/shows the Site Administration menu's Staging app.

Notice that some of the permissions above are related to the export/import functionality. Since these permissions are directly affected by the Publish Staging permission, they are important to note. Visit the Importing/Exporting Pages and Content section for more details on importing/exporting Site and page content.

CHAPTER 39

SCHEDULING WEB CONTENT PUBLICATION

Liferay's WEM lets you define when your content goes live. You can determine when the content is displayed, expired, and/or reviewed. This is an excellent way to keep your Site current and free from outdated (and perhaps incorrect) information. The scheduler is built right into the form your users access to add web content. Specifically, it can be found in the bottom panel listed with several other configurable settings.

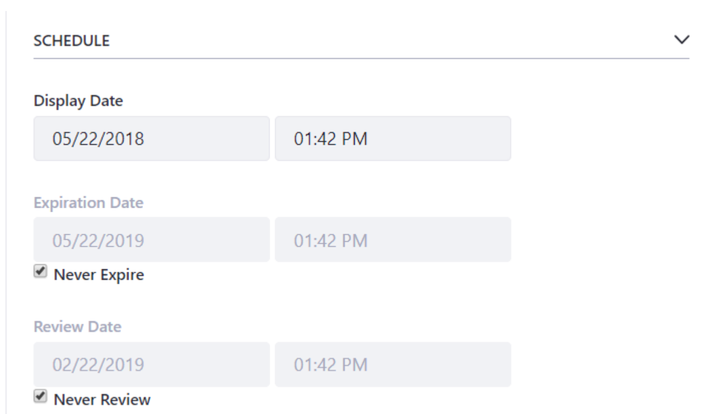

Figure 39.1: The web content scheduler can be easily accessed from the right panel of the page.

The scheduler offers several configurable options:
Display Date: Sets (within a minute) when content will be displayed.
Expiration Date: Sets a date to expire the content. The default is one year.
Never Expire: Sets your content to never expire.
Review Date: Sets a content review date.
Never Review: Sets the content to never be reviewed.
As an example, you'll step through the process of scheduling a web content article.

1. Navigate to the Product Menu → *Content* → *Web Content*.

2. Create a new web content article by selecting the *Add Web Content* button () → *Basic Web Content*.

3. Add content for your web content article.
4. Select the *Schedule* dropdown menu on the web content form. Configure the publication schedule.
5. Click *Publish*. Your web content article is now created and abides by the scheduling parameters you've set.

When you set a Display Date for an existing article it does not affect previous versions of the article. If a previous version is published, it remains the same until the new version is scheduled to display. However, the expiration date affects all versions of the article. Once an article has expired, no version of that article appears.

Tip: If you want only the latest version of articles to expire, and not every past version, go to *Control Panel* → *Configuration* → *System Settings* → *Web Content* → *Virtual Instance Scope* → *Web Content* and uncheck *Expire All Article Versions Enabled*. This makes the previously approved version of an article appear if the latest version expires.

The scheduling feature gives you great control in managing when, and for how long, your web content is displayed on your Site. Additionally, you can determine when your content should be reviewed for accuracy and/or relevance. This makes it possible to manage your growing inventory of content.

Chapter 40

Restoring Deleted Assets

Have you ever had that life-altering experience where you deleted an important file and immediately regretted deleting it? The deed is usually followed by a palm to the forehead or a sick feeling. Good news! Liferay DXP is here to turn that frown upside down with the *Recycle Bin* feature. With the Recycle Bin, the *Move to the Recycle Bin* action replaces *Delete* for certain asset types. Content is now temporarily stored in the Recycle Bin. This allows the content to be restored back to its original state. Recycled items can expire after a certain period of time, resulting in their permanent deletion. Before diving into how the Recycle Bin works, you'll look at how to configure it.

40.1 Configuring the Recycle Bin

The Recycle Bin supports instance-wide scope or site-specific scope. The instance-wide scope of the Recycle Bin is set by adding the `trash.enabled` property to your `portal-ext.properties` file. By default, the Recycle Bin is enabled instance-wide. You'll go into more detail for adding this property and several others to your properties file later in the section. First, you'll explore the UI and see what the Recycle Bin can do.

First, you'll configure the Recycle Bin for site-specific scoping.

1. Choose the Site you'd like configure for the Recycle Bin from the Site Administration menu.

2. Click *Configuration → Site Settings*.

3. Next, select the top *Advanced* tab and click *Recycle Bin*. You'll notice a few configurable options.

   Enable Recycle Bin: enable and disable settings for the Recycle Bin's site-specific scope.

   Trash Entries Max Age: customize the number of minutes a file is kept in the Recycle Bin until its permanent deletion (default is 43200 minutes, or 30 days).

4. When you've finished configuring your Recycle Bin settings, click *Save*.

Note: If you disable the Recycle Bin while it's still holding recycled items, the recycled items remain stored and reappear in the Recycle Bin if it is re-enabled.

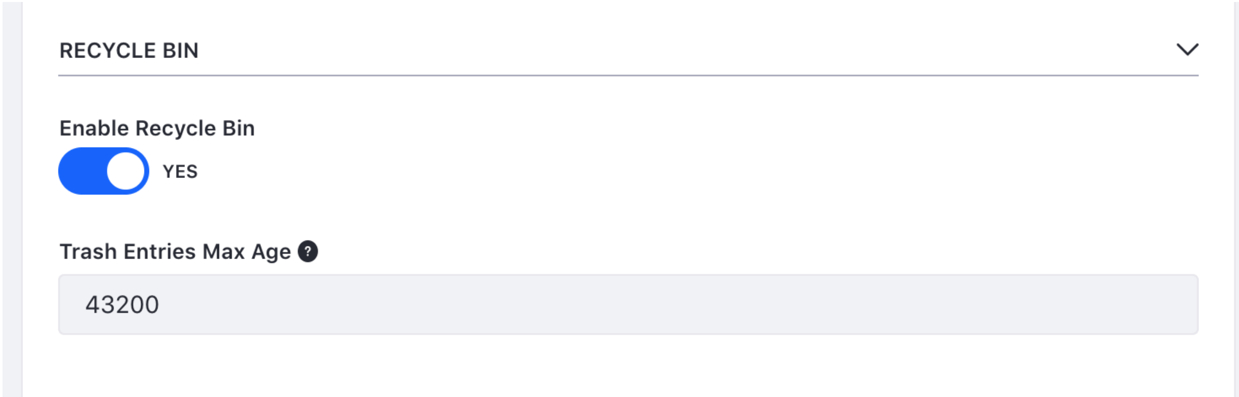

Figure 40.1: The Recycle Bin offers several configurable options for your site.

You can also configure the Recycle Bin via properties in the `portal.properties` file. Remember that it's a best practice not to edit the `portal.properties` directly, but to create a separate `portal-ext.properties` file containing the properties to override. There are some additional options not available in the GUI that you can set:

`trash.search.limit=500`: set the limit for results used when performing searches in the Recycle Bin (default is 500).

`trash.entry.check.interval=60`: set the interval in minutes for how often the trash handler runs to delete trash entries that have been in the Recycle Bin longer than the maximum age (default is 60).

Also, as was mentioned earlier, there are properties to enable the Recycle bin instance-wide and set trash entries' maximum age.

`trash.enabled=true`: set this property to *false* to disable the Recycle Bin for all sites in the portal (default is *true*).

`trash.entries.max.age=43200`: set the number of minutes trash entries should be held before being permanently deleted.

Visit the portal.properties file to view all of the configurable properties for the Recycle Bin.

Next, you should make sure permissions are set properly for users who can handle/view the assets in the Recycle Bin. Users who had *View* permissions on a document when it was recycled can also view that document in the Recycle Bin. Users who had *Update* or *Delete* permissions on a document when it was recycled can restore the document.

Now that you've successfully configured the Recycle Bin, you'll look at how to use it.

40.2 Using the Recycle Bin

The Recycle Bin is temporary storage where assets go when you delete them. You can recycle several different types of assets:

- Blogs
- Bookmarks
- Documents and Media
- Message Boards (and attachments)

- Web Content
- Wiki (and attachments)

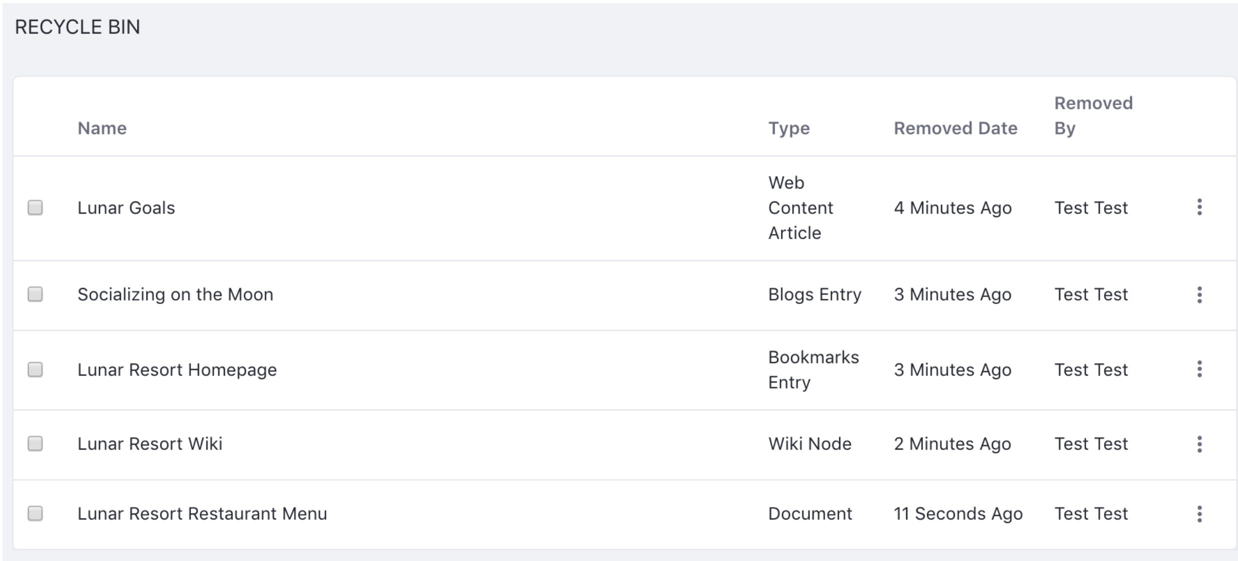

Figure 40.2: The Recycle Bin provides a seamless administrative experience for deleting and removing content.

Note: Attachments added to Wiki and Message Board entries do not go to the Recycle Bin when they are deleted. They can be restored in a similar fashion from the *Removed Attachments* menu within the application.

To demonstrate using the Recycle Bin let's delete a web content article and then restore it. You'll run through two different methods of restoring the file.

1. Go to Site Administration and select *Content* → *Web Content*.
2. Select the *Add* button (+) and click *Basic Web Content*.
3. Enter some text for the Title and Content and click *Publish*.
4. Click the article's *Actions* button (⋮) and click *Move to the Recycle Bin*.

   Note that the *Delete* button is not listed. Liferay DXP avoids the risk of accidental deletion of your files by funneling the content through the Recycle Bin.
5. After deleting the file, a success message appears, offering an *Undo* option. Click *Undo*. The web content is retrieved from the Recycle Bin and stored in its original place.
6. Click *Move to the Recycle Bin* again.
7. Go back to Site Administration and click Recycle Bin from the Content dropdown.
8. Find your sample web content and click its *Actions* button.

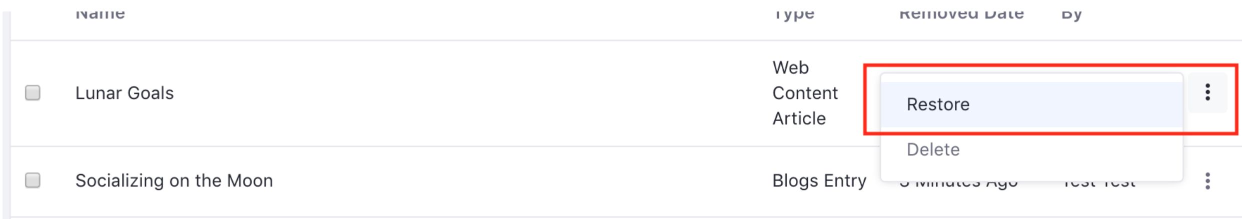

Figure 40.3: In the Recycle Bin, you have the option of restoring or permanently deleting the content.

9. You can restore or delete the content. Select *Restore*.

10. Navigate back to the Web Content screen and notice that your sample web content is back to its original place.

That covers the two general processes of sending and restoring content to/from the Recycle Bin. For other asset types, the Recycle Bin works similarly.

Some applications, such as Web Content and Documents and Media, support folders for organizing content. You can also send folders to the Recycle Bin. Keep in mind that this sends any sub-folders of the deleted folder all the files it contains to the Recycle Bin. Folders are restored and deleted the same way as a single file.

Delete within the Recycle Bin is the permanent delete button. Once you select this, your file cannot be retrieved and is gone forever. There is also an *Empty the Recycle Bin* option accessible from the () button at the top of the Recycle Bin screen. This permanently deletes all the files from the Recycle Bin.

Drag and Drop

You can also drag and drop items into the Recycle Bin. While you're in the Control Panel, select an asset and drag it to the Recycle Bin portlet on the Control Panel menu. When you click and begin dragging the asset, a message appears near your cursor notifying you of the number of files ready to be moved, and the Recycle Bin is highlighted, showing you where the files can be dropped. After you drop the asset onto the Recycle Bin portlet, the asset is removed from its original location and transferred to the Recycle Bin.

40.3 Recycle Bin Intelligence and Support

Have you ever wondered what happens to file shortcuts if their linked assets are recycled? What if you restore a file that has the same name as another file currently stored in your site/instance? The Recycle Bin already knows how to handle these types of issues.

When documents with shortcuts are moved to the Recycle Bin, the shortcuts are removed. This ensures that all your links and shortcuts work and cuts down on maintenance time and backtracking.

Another important trait is the duplicate name recognition feature. When a file is restored, the Recycle Bin scans the corresponding asset type files currently in the site/instance to check for duplicate file names. If a duplicate file name is found, the Recycle Bin prompts you to overwrite the existing file or rename the file name you're trying to restore.

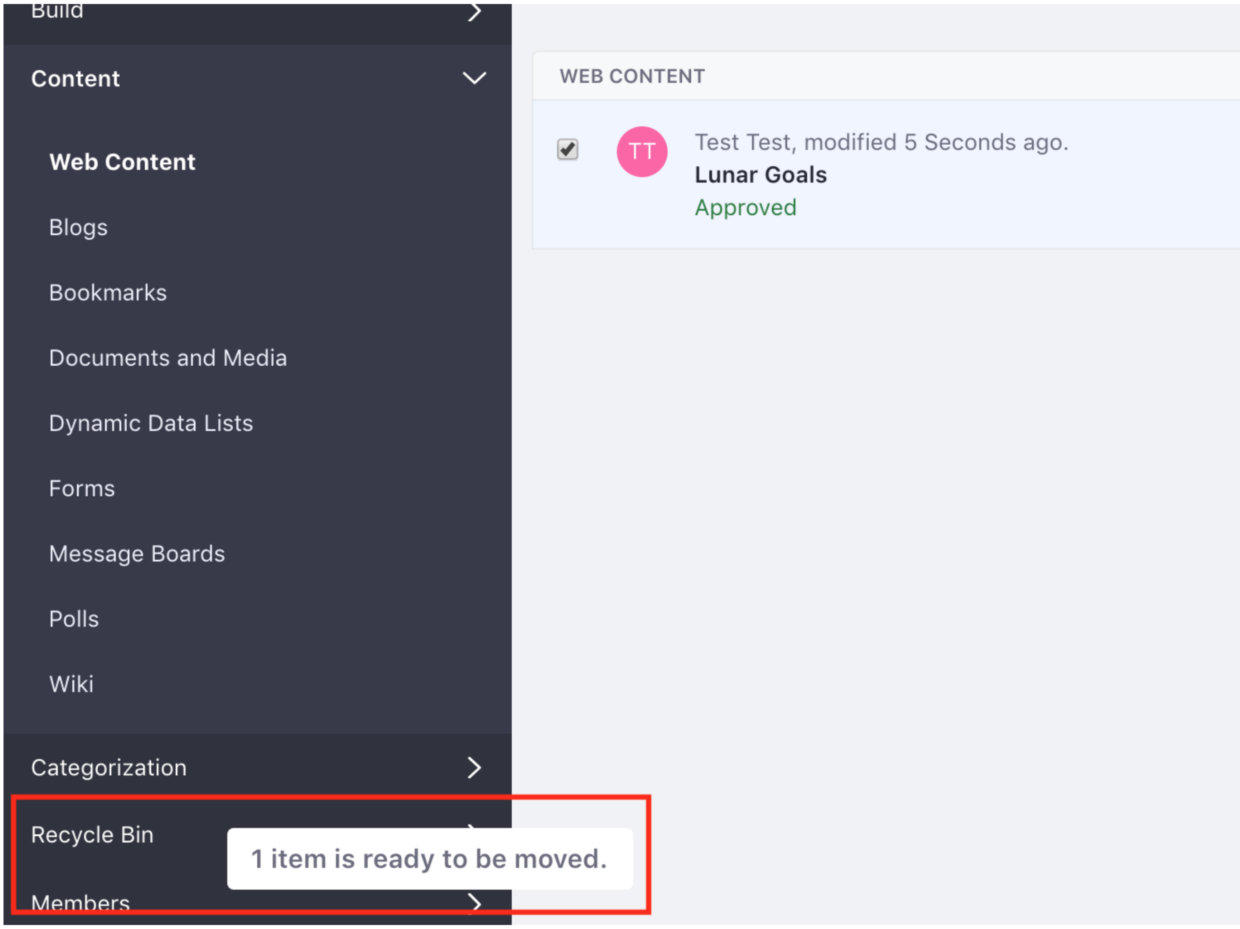

Figure 40.4: A quick and easy way of disposing your items is the drag and drop method.

For example, suppose you have the document `file1` stored in the Recycle Bin and you have a separate document you created later with the same name in the document library. If you try to restore the `file1` document, the Recycle Bin recognizes duplicate names and prompts you to overwrite the existing document in the document library or rename the document you're trying to restore.

Although the Recycle Bin prohibits the restoration of files that match pre-existing file names in your site/instance, it stores files with matching names.

The Staging Recycle Bin

Although you there is only one master Recycle Bin for all asset types, when staging is enabled a *Staging* Recycle Bin is created. The original Recycle Bin, or *Live* Recycle Bin, is still viewable while in staging; however, it is never used.

During staging, everything you recycle is sent to the Staging Recycle Bin. This prevents staged and unstaged recycled content from mixing. For example, if you have an unstaged document currently on your live site you can enable staging and delete that document. If you were to turn staging off and return to the live site, without separate Recycle Bins, the live document would be both on your site and in the Recycle Bin! Because of this, the separate Staging Recycle Bin is

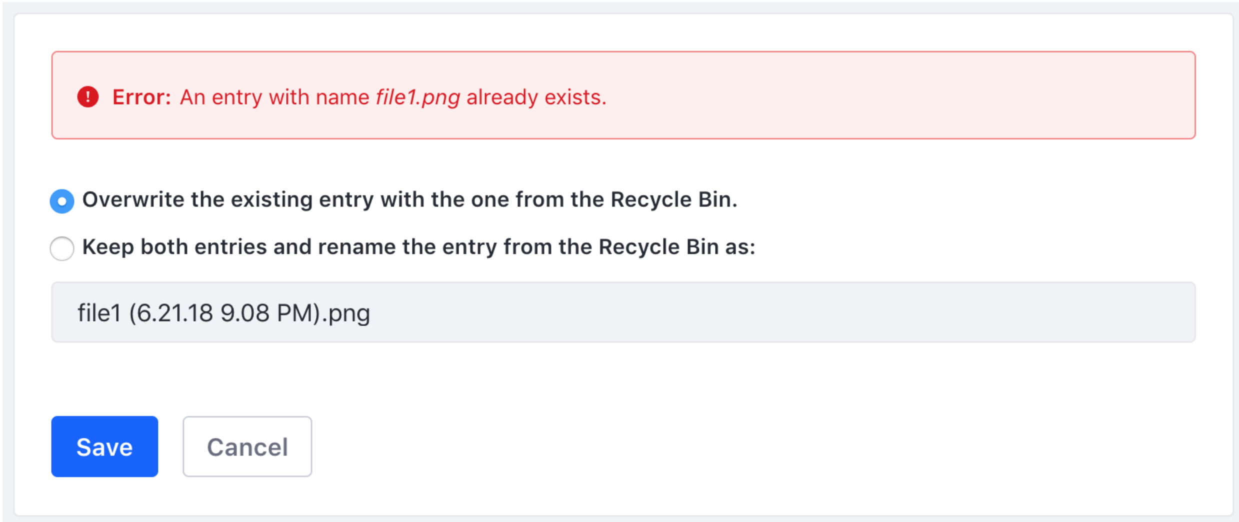

Figure 40.5: The Recycle Bin always scans your site/instance for duplicate file names during the restoration process.

necessary and only used during the staging process. When you publish your staged material, the Staging Recycle Bin content is transferred to the Live Recycle Bin.

Note: The Staging Recycle Bin saves its contents until the staged material has been published to the live site. This means that you can turn the staging mode on and off without losing your recycled material.

The Recycle Bin saves you time by letting you restore content that's been recycled. Instead of recreating or re-uploading content, you'll be tailoring your Liferay instance to fully leverage its capabilities.

Chapter 41

Collaboration

Liferay DXP contains an expansive collaboration suite that empowers users to create content and communities that they couldn't create alone. A robust document management system is a key component of this suite. As users produce digital assets—documents, videos, audio—they can store and share them using the Documents and Media Library. Documents and Media supports file check in and check out to prevent conflicting edits from multiple users, and maintains a version history of those files. It also contains its own repository, and for added flexibility can connect to external repositories. Once files exist in Documents and Media, users can insert them in other content like blog posts and wiki articles.

The collaboration suite also contains apps that let users share information and create active communities. The Message Boards app gives users a platform for discussions. The Blogs app lets users publish their ideas using rich content. Notifications keep users informed of what's happening. Social networking apps let users connect and share in ways that bolster friendship and productivity. And this is just scratching the surface—there are many more apps that help users communicate, produce, and present.

The guides that follow show you how to leverage these features, and more, in detail.

Chapter 42

Managing Documents and Media

The Documents and Media library stores files on the server using the same type of structure that you use to store files locally. It accepts files of any kind, can serve as a virtual shared drive, and can mount and browse external repositories. You can organize documents using customizable document types and metadata sets and display them with automatic document preview generation. Its companion app, the Media Gallery, displays selected content from the Documents and Media library. It can render image, audio, and video files.

Liferay Sync synchronizes Documents and Media folders with local folders on your devices, both your desktop machines and mobile devices.

You'll get started with Documents and Media by exploring how to publish files.

Chapter 43

Publishing Files

As you create sites, you'll probably want to share files on them. The Documents and Media library (Document Library) lets you upload and publish all kinds of files on your sites. Pictures, videos, spreadsheets, slide presentations, and more can be stored in and shared from the Document Library. Document Library instances can be scoped to a portal instance, site, or page, so you can work with files where they're relevant.

In these tutorials, you'll learn how to add files, display them, and collaborate on them. You'll learn how to use both the Documents and Media Library and the Media Gallery. And lastly, you'll learn how to collaborate on files from within several environments, including your browser and local desktop file system.

43.1 Adding Files to a Document Library

This article covers the following topics to help you get started adding files to your Document Library:

1. **Granting File Permissions and Roles:** Determine who can add, view, and update files. Doing this before adding files ensures that only those you wish can access your Document Library.

2. **Adding Files:** Add specific types of files and their associated metadata to your Document Library.

Granting File Permissions and Roles

You should carefully manage who can add, view, and update files. You can store files of all kinds for various purposes. For example, you may have one set of files intended for only specific site members and another intended for everyone, including guests. You can use Roles and Permissions to control access to Document Library files. The Document Library's folder permissions also help you organize files.

Follow these steps to create a Role for managing files in your site's Documents and Media:

1. Open the *Menu* () and navigate to *Control Panel* → *Users* → *Roles*.

2. Select the *Site Roles* tab (or *Organization Roles*, for an Organization Role) and then click the *Add* button () to begin creating a role.

3. Give your Role a name and a description, then click *Save*.

4. Select your Role's *Define Permissions* tab. In the Role's permission definition screen, navigate to *Site Administration* → *Content* → *Documents and Media*. In the *General Permissions* section, select *Access in Site Administration* and click *Save*.

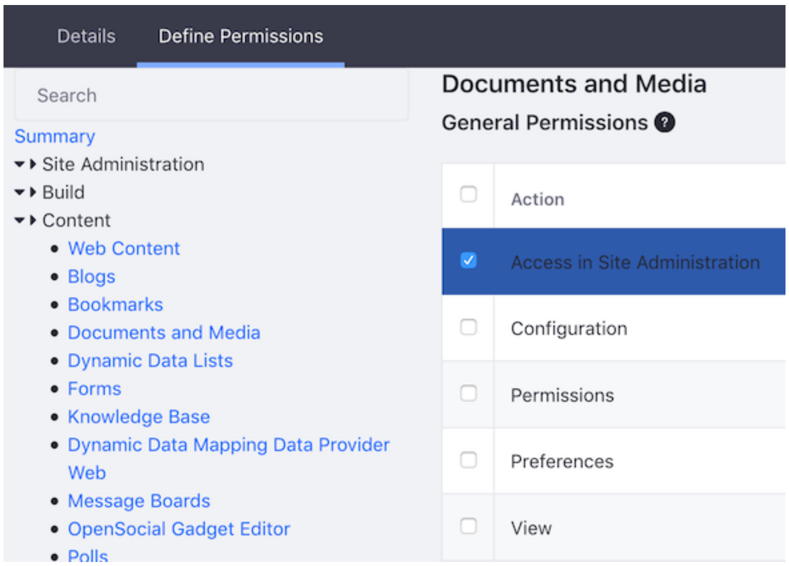

Figure 43.1: It's often helpful to define a role for specific users to access Documents and Media from Site Administration.

5. Assign this Role to the Users that should manage media. For more information on this and other topics related to Roles, see the documentation on Roles and permissions.

Using the Add Menu

Follow these steps to add files to your site's Document Library:

1. Open the *Menu* (), click on your site's name, and navigate to *Content* → *Documents and Media*. The Documents and Media screen appears and displays the Documents and Media library's *Home* (its root folder). As you add files and folders to the Document Library, they're listed here.

2. Click the *Add* icon () and select the type of document to add to the Document Library. You can add documents, folders, and shortcuts much like you would on a desktop file system. You can even configure access to an entirely different repository. The Add menu's options are described below.

3. When you're finished selecting the file to upload and filling out any document type fields that are necessary, click *Publish*.

Folder: Create a new folder in the app's file system.

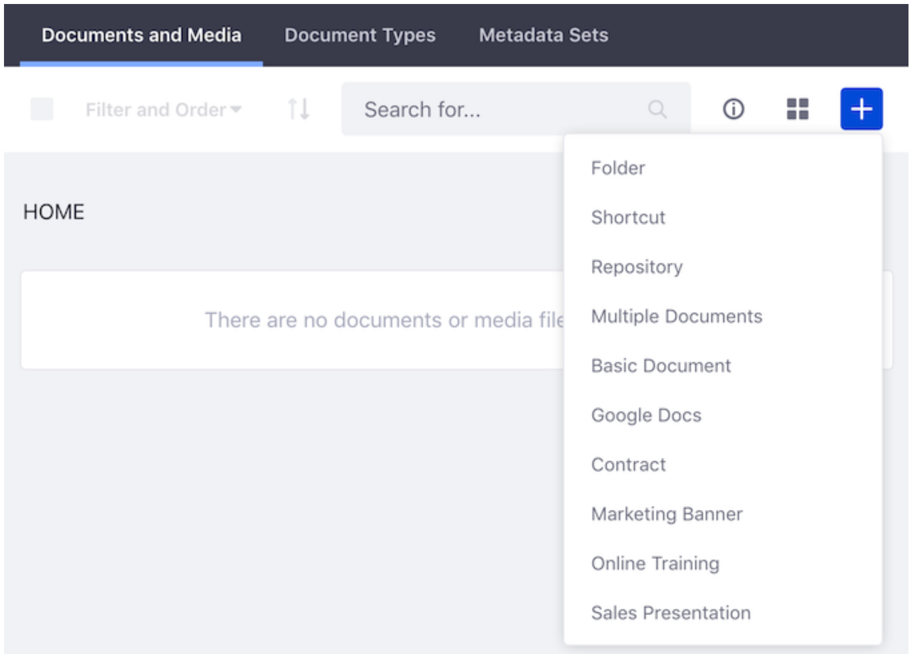

Figure 43.2: The Documents and Media's *Home* folder starts empty. But the Add menu lets you upload and add all kinds of documents to the library.

Shortcut: Create a shortcut to any document that you can view. You can set permissions on the shortcut to specify who can access the original document via the shortcut.

Repository: Add access to an external repository. See the store types documentation for more information.

Multiple Documents: Upload several files at once. You can apply a single description and document type to all the files. You can also categorize and tag the files, and assign them default permissions.

Basic Document: Upload a single file. By default, basic documents aren't described by any metadata sets.

The remaining items in the Add menu are default document types described by a unique metadata set. When you add a document belonging to a document type, a form appears that lets you pick the file to upload and enter the data defined by the document type's metadata set. For example, the *Contract* document type is used to describe legal contracts. By default, contracts are described by fields for the effective date, expiration date, contract type, status, legal reviewer, signing authority, and deal name.

Any custom document types that have been defined also appear in the Add menu.

43.2 Creating Folders

You'll need folders to organize all but the most limited set of files. Here, you'll learn how to work with folders in a Document Library:

- Adding a Folder
- Document Type Restrictions and Workflow

- Setting Folder Permissions

Adding a Folder

Follow these steps to add a folder:

1. Open the *Menu* ( ), click on your site's name, and navigate to *Content* → *Documents and Media* for your site. The Documents and Media screen appears and displays the Documents and Media library's *Home* (its root folder).

2. Click the *Add* icon ( ) and select *Folder*. The New Folder form appears.

3. In the New Folder form, name and describe your folder. Then expand the *Permissions* section.

4. In the Permissions section, set the folder's permissions. The *Viewable by* menu lets you select who has View permission for the folder:

    - Anyone (the Guest role; this is the default option)
    - Site Members
    - Owner

    Click the *More Options* link to choose the other folder permissions for the Guest and Site Member roles. By default, site members can add files, subfolders, shortcuts, and subscribe to changes to the folder's files. Guests don't have any such permissions, which is typically what you want.

    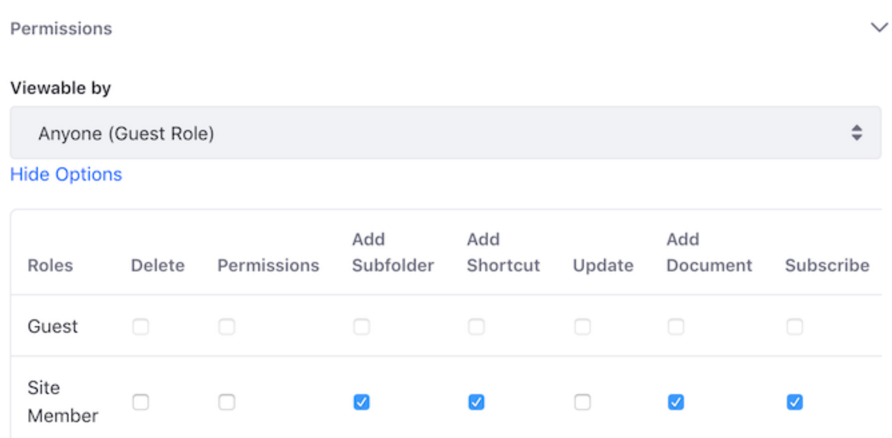

    Figure 43.3: Select your folder's permissions.

5. To finish creating the folder, click *Save* after making your selections in the Permissions section.

Upon creating the folder, it appears in your Document Library. Opening the folder's *Actions* menu ( ) presents several options for managing the folder. The following sections describe some of these options.

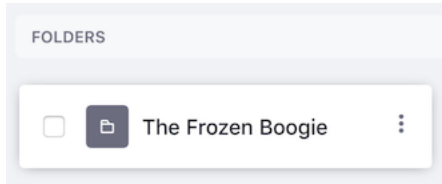

Figure 43.4: Your new folder appears in the Document Library.

Document Type Restrictions and Workflow

After creating a folder, you can restrict what document types are allowed in it. You can also choose what workflow (if any) to use for approving files added to or edited in the folder.

Follow these steps to change a folder's document type restrictions and workflow:

1. Click the folder's *Actions* menu ( ) and select *Edit*.
2. Expand the *Document Type Restrictions and Workflow* section. In this section, choose from the following options:
    - Use Document Type Restrictions and Workflow of the Parent Folder (the parent folder)
    - Define Specific Document Type Restrictions and Workflow for this Folder (the current folder)
    - Default Workflow for this Folder (the current folder)
3. Click *Save* when you're finished.

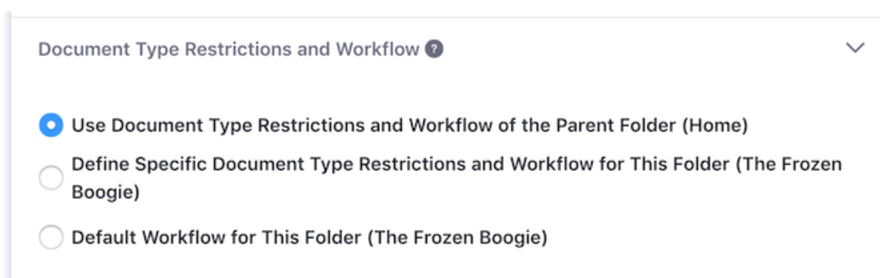

Figure 43.5: You can set the document type restrictions and workflow to use for a folder's files.

Setting Folder Permissions

When creating a folder, you can set some of its permissions via the new folder form. Fine tuning a folder's permissions, however, can only be done after creating the folder.

Follow these steps to fine tune a folder's permissions:

1. Click the folder's *Actions* menu ( ) and select *Permissions*. The Permissions window appears.
2. In the Permissions window, set the permissions you want to use for this folder. The permissions listed below are available for each role.

3. Click *Save* when you're finished setting permissions.

Access: Access the folder's contents.
Update: Edit the folder's attributes and/or move the folder under a new parent folder.
Add Subfolder: Create folders within the folder.
Add Shortcut: Create a shortcut (link) to any file in the folder that the role is authorized to view.
Subscribe: Receive email notifications when files are added to or modified in the folder. Note that you can specify the email sender and template from the Documents and Media's *Options* () → *Configuration* menu.
Add Document: Add a new file to the folder.
Permissions: View and modify the folder's permissions.
Delete: Move the folder to the Recycle Bin.
View: View the folder.

43.3 Using the Documents and Media Management Bar

The Documents and Media *Management Bar* is where people that manage documents go to unwind after a long day at work. Just kidding. The Management Bar, as its name implies, contains tools for managing the files and folders in your Document Library. It appears above the files and folders in Documents and Media.

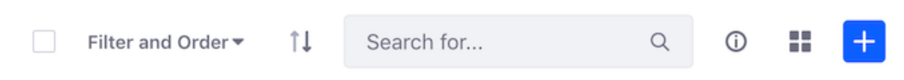

Figure 43.6: The Management Bar is a great place to hang out if you're managing documents.

If you've added files or folders to your Document Library, then you're already familiar with the Management Bar's *Add* button (). The sections that follow describe the rest of the Management Bar.

Note: If a Document Library contains more items than it can display at once, you can use the navigation tool that appears at the bottom of the window to switch your view to another page or configure the page to display more items per page.

View Types

The *View Types* button is to the left of the Add button. It lets you choose how to display the Document Library's items. The View Types button's icon depends on the selected view type:

- **Cards** (): Shows a card-like rendering of the item. If the item isn't an image, a generic image for the item's type is displayed. For files, each card also contains the file's suffix (e.g., JPG, PNG, etc.), timestamp, name, and workflow status (e.g., Approved, Draft, etc.).

- **List** (): Shows the same information as the Cards view type, in a list with small file renderings.

- **Table** (⊞): Shows the same information as the other view types, in a list with no file renderings. Also, the file information is in columns.

The items in all view types have an Actions menu (⋮). These actions are also available in when viewing each item separately.

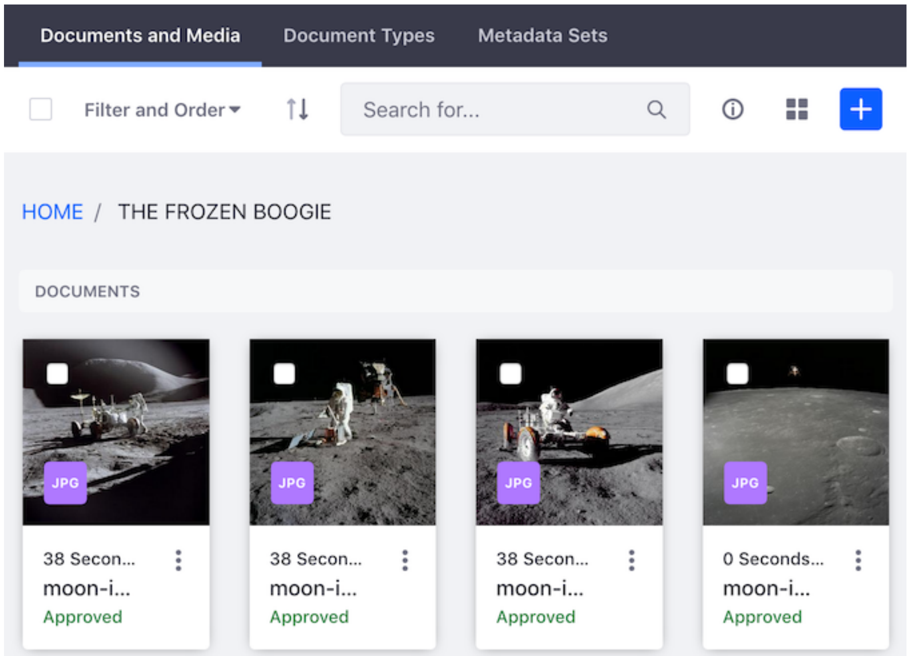

Figure 43.7: The Cards View type shows items in large card-like renderings.

The Info Panel

To display an info panel with the current folder's details, click the *Information* icon (ⓘ). The info panel slides out from the right side of the screen and contains the folder's name and number of items. It also has these buttons:

- **Subscribe** (☆): Get notifications about files added to or modified in the folder.
- **Actions** (⋮): Lists actions you can perform on the current folder.

Finding and Arranging Items

The Management Bar also contains tools that help you locate and arrange items in the Document Library. The most prominent of these tools is the *Search* bar, where you can find files by keywords.

To the left of the Search bar, the Sort button (↑↓) arranges items in ascending or descending order.

You can also arrange items via the *Filter and Order* selector using these criteria:

- **All:** Shows all of the current folder's immediate subfolders and files (default).
- **Recent:** Shows the most recently modified files.
- **Mine:** Shows all the current user's files (no matter their folder).

- **Document Types:** Shows the files of the selected document type. Upon choosing this option, you must select the document type you want from a popup.

You can also select from the following criteria for ordering items:

- Size
- Downloads
- Modified Date (default)
- Create Date
- Title

Selecting Items

The checkbox on the left-most side of the Management Bar selects all currently displayed items. Selecting multiple items lets you act on all of them at once. You can also select multiple items individually by using the checkboxes for each. When you select one or more items, the Management Bar changes to reflect the actions you can take on the selected items.

Figure 43.8: With items selected, the Management Bar changes.

Here are the actions you can take on the selected items:

- Download ()
- Checkin ()
- Checkout ()
- Move ()
- Move to Recycle Bin ()

File checkout and checkin is explained in Checking out and Editing Files.

43.4 Viewing File Previews

File previews help users browse and find media efficiently. To view a preview of a file, click the file's name in the Document Library. If the file is an image, the image appears. If an app is installed that can render a preview of the file type, a representative image of the file appears (e.g., the opening frame of a video file or a presentation's first slide). If there are no such preview apps for the file, a generic image based on the file type appears.

File Preview Apps

Whenever possible, Liferay DXP generates previews of documents added to the Document Library. Out of the box, Java-based APIs generate previews. The only tool available that is 100% Java and has a compatible license to be distributed with Liferay DXP is PDFBox. A separate thread generates a preview for PDFs when uploaded. This process may last only a few seconds for a small file. The larger the file, the longer it takes.

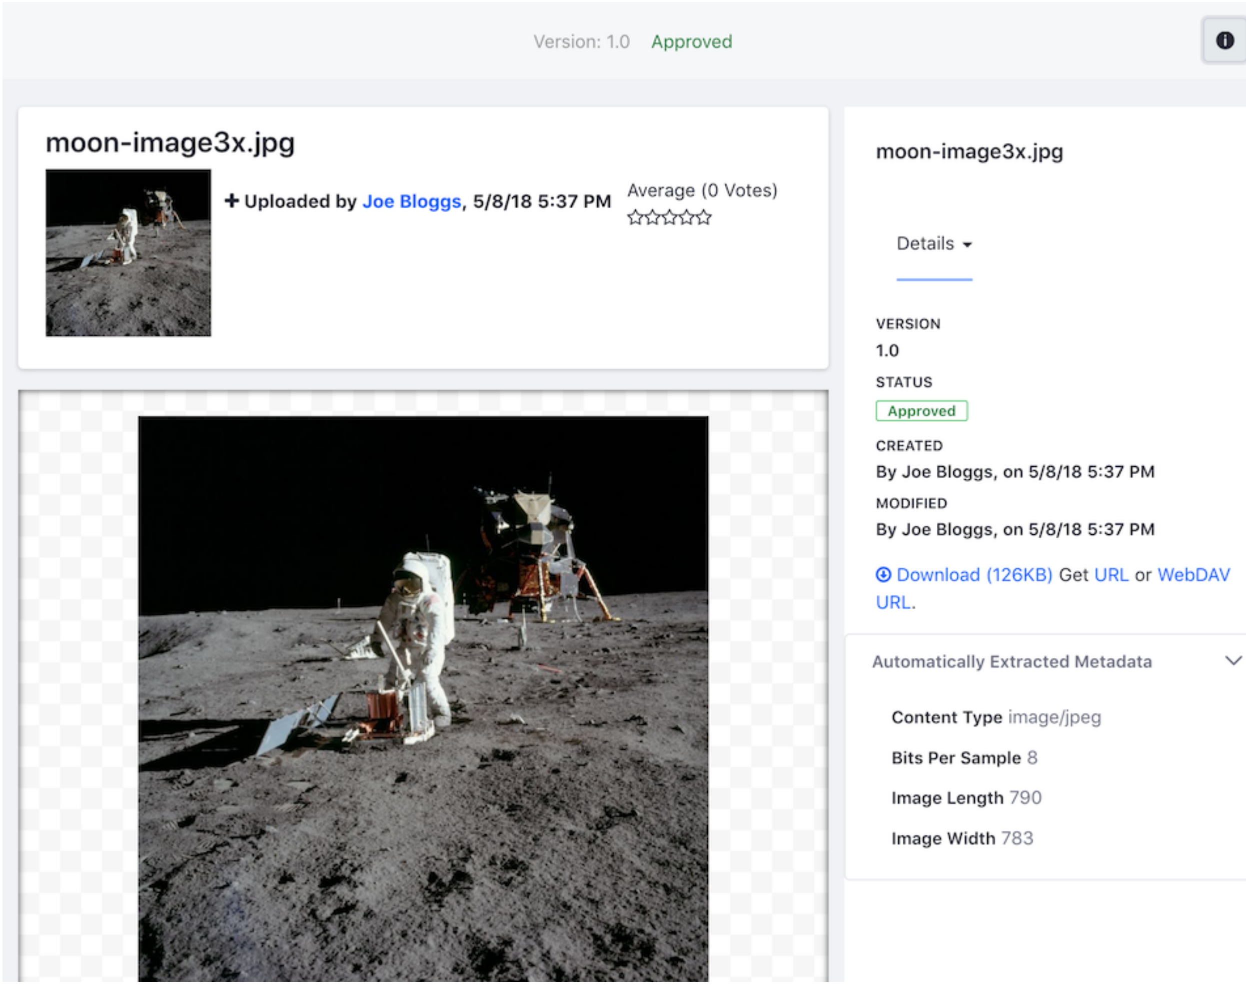

Figure 43.9: File previews let you view and manage a file.

While PDFBox provides a default implementation of image generation for document previews and thumbnails, you must install and configure additional tools to harness the full power of document previews. These tools include:

- OpenOffice or LibreOffice: Using one of these in server mode lets you generate thumbnails and previews for supported file types (`.pdf`, `.docx`, `.odt`, `.ppt`, `.odp`, etc.), view documents in your browser, and convert documents.
- ImageMagick (also requires Ghostscript): Enables faster and higher-quality previews and conversions.
- Xuggler: Enables audio and video previews, lets you play audio and video files in your browser, and extracts thumbnails from video files.

After installing these tools, you can configure them via portal properties in the Control Panel's Server Administration screen, or in a `portal-ext.properties` file. To learn how to use these tools, refer to Configuring Liferay DXP.

With these tools installed and configured, a customized viewer displays Documents and Media content, depending on the content type. For example, you can view a document with a customized viewer that lets you navigate through the document's pages. You can also view and play multimedia documents (audio or video). If the browser supports HTML5, the viewer uses the browser's native player. Otherwise it falls back to a Flash player.

Managing Files

You can also manage a file from its preview. The file's icon, name, author, upload timestamp, and rating appear above the preview area. The comments area (below the preview area) lets you comment on and subscribe to comments on the file.

You can access more options from the *Options* menu (⋮) at the top-right of the screen. Here are the file options:

Download: Download the file.

Edit: Modify the file's name, description, document type, categorization, and related assets. You can even upload a new file to replace it. Note that modifying the file increments its version.

Edit With Image Editor: Edit the image in the Image Editor. The Image Editor is explained in Editing Images.

Move: Relocate the file to a different parent folder.

Checkout/Checkin: Checkout prevents others from editing the document while you are working on it. Other users can still view the current version of the document, if they have permission. You can check in the document when you're done with it.

Permissions: Specify which actions each role can perform on the file.

Move to Recycle Bin: Move the file from the Documents and Media library to the Recycle Bin.

Click the *Info* icon ( ) to view information about the file in an info panel that slides out from the right. A selector menu at the top of the info panel is set to *Details* by default. This shows the following information:

Version: The file's version.

Status: The file's workflow status (e.g., Approved).

Created: The user that created the file, and when it was created.

Modified: The user that last modified the file, and when it was last modified.

Download: A link to download the file.

URL: A URL to access the file in the Document Library.

WebDAV URL: A WebDAV URL for accessing the file via a desktop.

Automatically Extracted Metadata: Any and all metadata automatically extracted from the file. When adding new documents or viewing existing documents, a process is triggered automatically that extracts the file's metadata. The library used by this process is TIKA and it's included out of the box. Depending on your file's type and the metadata written with the file, you can find out all kinds of details. In the case of audio or video files, the media's duration is displayed.

To instead view the file's version history, select *Versions* from the selector menu. The info panel then changes to list the different versions of the file and lets you view, download, remove, and revert to specific file versions. File version history actions are explained in the guide Collaborating on Files.

43.5 Editing Images

Editing and re-uploading images when you only need to apply simple edits is tedious. Docs & Media contains a simple built-in image editor for exactly this reason. To access the image editor, locate the image you want to edit. Click the Actions icon (⋮) and select *Edit With Image Editor*.

You can also access the image editor when selecting an image to insert in content (i.e., via an item selector). Anywhere an image is, you can edit it. For example, you can access the image editor via item selector preview windows in blog entries and web content articles. To do this, click the pencil icon ( ) in the bottom-right corner of the preview window.

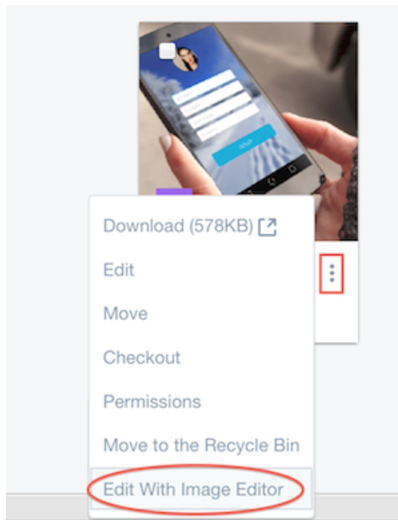

Figure 43.10: You can access the image editor through the Documents and Media repository.

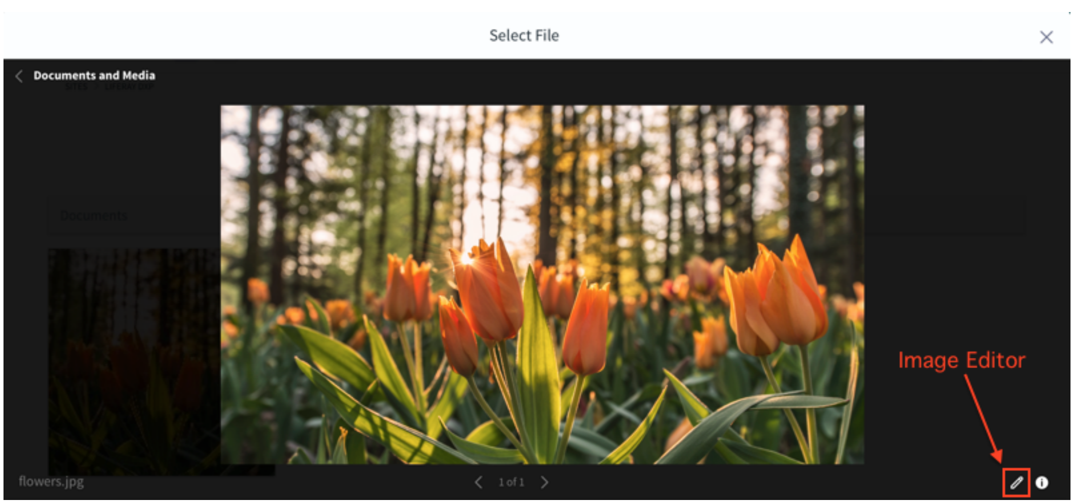

Figure 43.11: You can also access the image editor through the item selector preview window.

If you edit and save the image via the Documents and Media repository, the file version is incremented a minor version (e.g., from version 1.0 to version 1.1). You can view the image's version history (and previous versions) by clicking the image, clicking its *Info* button ( ), and then selecting *Versions* from the menu currently set to *Details*. In contrast, if you edit and save an image via an item selector, a copy of the image is created and saved to the document library.

Liferay designed the image editor with quick editing in mind. It offers a minimal, user-friendly UI. The main toolbar consists of three buttons, each of which contains a subset of options:

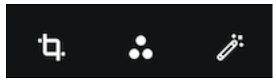

Figure 43.12: The image editor's UI is clear and to the point, offering only what you need.

Transform ()

- **Rotate**: Rotate the image to the left or right, in 90 degree increments.
- **Resize**: Resize the image. If the lock is closed, the aspect ratio is locked and changing width or height automatically adjusts the other dimension to maintain the aspect ratio. When the lock is opened, the width and height can be changed individually, letting the aspect ratio change (this isn't recommended because the image can become distorted).
- **Crop**: Crop the image.

Adjustment ()

- **Saturation**: Adjust the color saturation. The default value is 50. Values range from 0 (completely desaturated) to 100 (completely saturated).
- **Contrast**: Adjust the contrast. The default value is 50. Values range from 0 (no contrast) to 100 (full contrast).
- **Brightness**: Adjust the brightness. The default value is 50. Values range from 0 (completely black) to 100 (completely white).

Filter (): Apply a filter to the image.

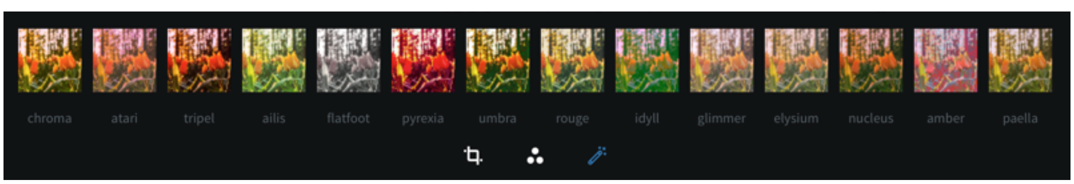

Figure 43.13: Select from a set of preset image filters.

Upon editing the image in the editor, you can click the *Cancel* button to cancel the changes, or the *Apply* button to apply them. Upon applying the changes, the history bar appears. It lets you undo, redo, or reset the changes. Use the Reset button with caution; it resets the image to its original state, reverting all changes made in the editor.

43.6 Publishing Files

Once your Document Library contains files, you may want to publish them in your site. Here are some ways to publish files:

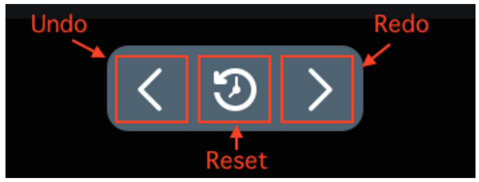

Figure 43.14: The history bar lets you undo, redo, and reset changes.

- Show them in a Documents and Media app.
- Display them in a Media Gallery.
- Use the Asset Publisher.
- Insert them in an asset like a web content article or blog entry.

Here, you'll learn to use the Media Gallery.

Using the Media Gallery

The Media Gallery publishes your media files in a simple gallery-like style. It shows a large thumbnail of each media file, lets the user download files, and has slideshow capabilities. A common way to use the Media Gallery is to create a separate page for displaying media and add a Media Gallery widget to it. This way, your media takes center stage.

Follow these steps to create a page that contains a Media Gallery widget:

1. Create a page and navigate to it in your site.

2. At the top-right of the screen, click the *Add* icon () then navigate to *Widgets* → *Content Management* and select *Add* next to *Media Gallery* (alternatively, drag the Media Gallery onto your page). The Media Gallery widget appears on the page.

3. Configure the Media Gallery widget to show your files. By default, it shows files from the Home folder of your site's Documents Library. To choose a different folder, click the widget's Options icon ( ) and select *Configuration*.

   The Configuration window appears and shows the *Setup* tab. This tab contains these sections:

   Display Settings: Lets you show each file's actions, filter the media types to display, and choose a display template for your media.

   Folders Listing: Lets you select a Document Library folder to serve as the root folder from which to display files. The root folder you select becomes the highest-level folder the Media Gallery can access. For example, if you create a subfolder of a parent folder, and then set that subfolder as the Media Gallery's root folder, the Media Gallery can no longer access the parent folder.

```
**Note**: To access the Carousel display template in Media Gallery, your
role must have View access for that template. Since the Carousel template is
in the Global scope, a Global-scope administrator must grant the role
permission to view the template.
```

```
![ You can configure the Media Gallery to use any Documents and Media folder as its root folder.](./images/dm-select-root-folder.png)
```

4. Configure the rest of the settings as desired in the Media Gallery app's other configuration tabs:

   Communication: Lists public render parameters the widget publishes to other widgets on the page. Other widgets can take action on these parameters. For each shared parameter, you can specify whether to allow communication via the parameter and select which incoming parameter can can populate it.

   Sharing: Lets you embed the widget instance as a widget on on any website, Facebook, Netvibes, or as an OpenSocial Gadget.

   Scope: Lets you specify the Document Library instance the widget uses: the current site's instance (default), the global instance, or the page's instance. If the page doesn't already have an instance of the widget, you can select *Your Page (Create New)* to create a page-scoped instance for the widget to display.

5. Click *Save* when you're finished configuring the Media Gallery widget.

The Media Gallery now shows your files, with images appearing as thumbnails. When you click a thumbnail, a slideshow appears showing the selected image. Below that image, thumbnails of the folder's other images are displayed. The slideshow continues until you click pause or view the last image. Closing the slideshow window returns you to the page.

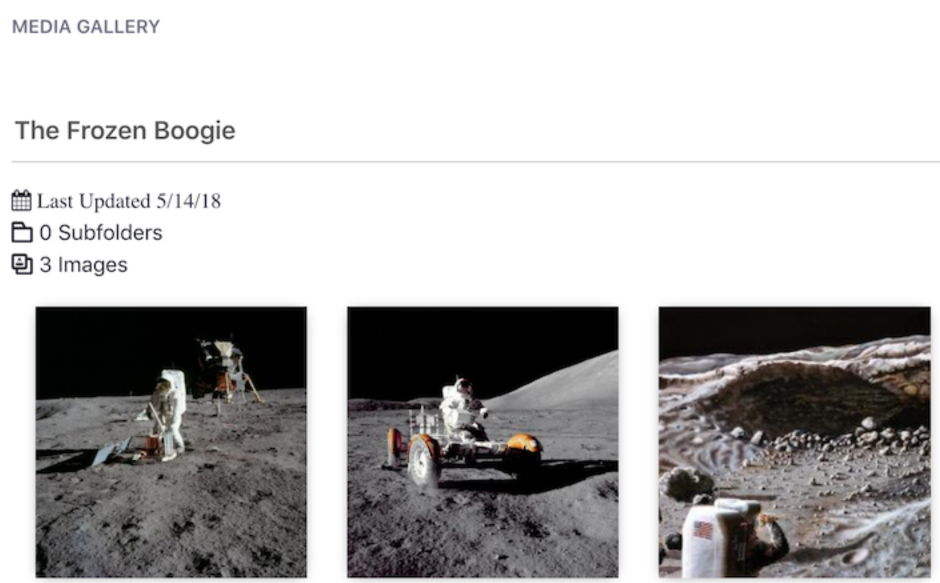

Figure 43.15: The Media Gallery renders large thumbnail images of media files.

43.7 Checking Out and Editing Files

When you check out a document in the Document Library, only you can make changes to it until you check it back in. This prevents conflicting edits on the same document by multiple users.

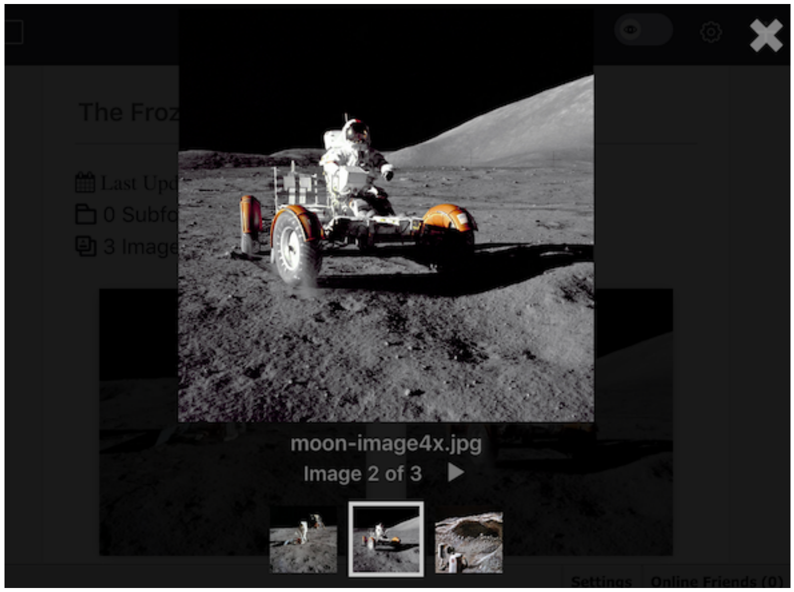

Figure 43.16: The Media Gallery's slideshow provides a nice way to view images.

When you check out a file, you can download it, replace it, move it to another Document Library folder, check it in, or cancel the checkout. Checking in a file also increments its version, which lets you keep track of changes.

Unless you're using Liferay Sync or a local drive mapped to the file's WebDAV URL, follow these steps to edit a Document Library file from your machine:

1. Checkout the file by clicking its Actions icon ( ) → *Checkout*. Upon checkout, the file's status changes to Draft.

2. Download the file by clicking its Actions icon ( ) → *Download*.

3. Edit the file locally.

4. Return to the Documents and Media Library and click the file's Actions icon ( ) → *Edit*. The file's edit screen appears.

5. From the file's Edit screen, select the edited local file for upload.

6. Click *Save and Check In*. In the pop-up that appears, select whether your change is a major or minor version, add any version notes that you need, and click *Save*.

Note: If you edit a file without checking it out, the file's edit screen displays a toggle for *Customize the Version Number Increment and Describe My Changes*. Setting this to *YES* lets you specify the version increment's type and description.

Follow these steps to access a file's version history:

1. Click the file in the Documents and Media Library.

2. Click the file's *Info* button ( ) at the top-right of the screen.

3. Select *Versions* from the menu currently set to display the file's details.

Each file version has an Actions menu ( ) that you can use to perform the following actions on that file version:

Download: Download the selected version of the file to your machine.

View: View the file entry screen for the selected version of the file.

Revert: Restores the selected file version as a new major file version.

Delete Version: Remove the file version from the Document Library. All other file versions remain intact.

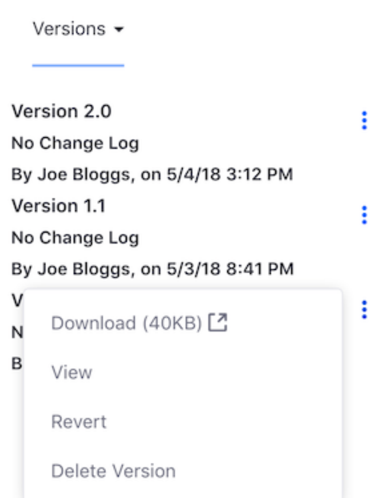

Figure 43.17: The version history actions let you inspect, delete, and reinstate file versions.

43.8 Desktop Access to Documents and Media

You can access the Document Library from your desktop file manager via WebDAV. WebDAV is a set of methods based on HTTP that let users create, edit, move, or delete files stored on web servers. WebDAV is supported by most major operating systems and desktop environments, including Linux, macOS, and Windows. Using your file manager via WebDAV doesn't bypass the functionality of the web interface—Liferay DXP increments the version numbers of files edited and uploaded via WebDAV.

To access the Document Library folder from a file browser, you must use your log-in credentials and the WebDAV URL of the folder you want to access. Follow these steps:

1. Navigate to the Documents and Media app that contains the folder you want to access. Click on the folder's Actions icon and select *Access from Desktop*.

2. Copy the WebDAV URL and follow the instructions for your operating system:

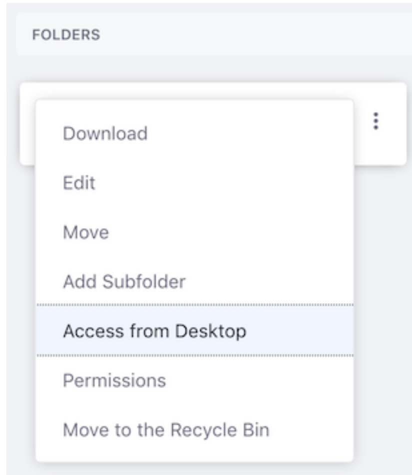

Figure 43.18: Select *Access from Desktop* to get the folder's WebDAV URL.

- **Windows:** Map a network drive drive to the WebDAV URL. Enter your credentials when prompted. The Document Library folder appears in the network drive. From your file browser, you can now add, edit, move, or delete files in this folder.
- **macOS:** In the Finder, select *Go* → *Connect to Server*. In the Server Address field, enter the WebDAV URL of the folder you want to access, then click *Connect* and enter your credentials when prompted.
- **Linux:** In your file manager, you must slightly modify the Document Library folder's WebDAV URL. For KDE's Dolphin, change the URL's protocol to `webdav://` instead of `http://`. For GNOME's Nautilus, change the URL's protocol to `dav://` instead of `http://`. Then press *Enter* and enter your credentials when prompted.

Now you can access the Document Library folder from your desktop file system. If you edit a file in this folder on your file system, the change also shows up in the same Document Library folder in the portal. What's more, the file's minor version is incremented due to the edit.

43.9 Linking to Google Drive™

You can create Document Library files that link to files in Google Drive™ and Google Photos™. This lets you access your Google files from the Document Library. Note that this functionality isn't available by default. To enable it, you must complete these steps:

1. Install the Liferay Plugin for Google Drive™ from Liferay Marketplace.
2. Create and/or configure a Google project capable of communicating with your Liferay DXP instance. The Google Picker API must be enabled for this project. This API lets you select Google files to link to. You must also create the credentials the Google project needs to communicate with your Liferay DXP instance.
3. Configure your Liferay DXP instance to communicate with your Google project.

This article shows you how to complete these steps and finishes with an example of linking to a Google file from the Document Library.

Important: The Liferay Plugin for Google Drive™ is a Labs application available for Liferay CE Portal and Liferay DXP. Labs apps are experimental and not supported by Liferay. They're released to accelerate the availability of useful and cutting-edge features. This status may change without notice. Use Labs apps at your own discretion.

Install the App

First, you must install the the Liferay Plugin for Google Drive™ from Liferay Marketplace. This app is available via the following links for Liferay CE Portal and Liferay DXP:

- Liferay Plugin for Google Drive - CE
- Liferay Plugin for Google Drive - DXP

If you need help installing apps from Marketplace, see the documentation on using Marketplace.

Configure Your Google Project

Follow these steps to create and/or configure your Google project so it can communicate with your Liferay DXP instance:

1. Go to the Google API Console. If you don't have a suitable project, create a new one.

2. Enable the Google Picker API for your project. For instructions, see the Google API Console documentation on enabling and disabling APIs.

3. Create an OAuth 2 client ID in your Google project. For instructions, see the Google API Console documentation on setting up OAuth 2.0. Enter these values when creating your client ID:

   - **Application type:** Web application
   - **Name:** Google Docs Hook
   - **Authorized JavaScript origins**: `[liferay-instance-URL]` (e.g., `http://localhost:8080` is the default for local development machines)
   - **Authorized redirect URIs**: `[liferay-instance-URL]/oath2callback`

4. Create a new API key in your Google project. For instructions, see the Google API Console documentation on creating API keys. Be sure to restrict the key to HTTP referrers (web sites), and set it to accept requests from your Liferay DXP instance's URL.

Your new OAuth client ID and public API access key now appear on your Google project's Credentials screen. Keep this screen open to reference these values as you specify them in Liferay DXP.

Configure Liferay DXP's Google Apps Settings

Configure your Liferay DXP instance with the client ID and API key that you created in the previous step:

1. Navigate to *Control Panel* → *Configuration* → *Instance Settings*.
2. Click the *Miscellaneous* tab and expand the *Google Apps* section.
3. For *Google Apps API Key*, enter the Google API key that you created in the previous section.
4. For *Google Client ID*, enter the Google OAuth client ID that you created in the previous section.
5. *Save* your changes.

Creating Linked Files

With the preceding configuration steps complete, you can create files in your Document Library that link to files in Google Drive™ or images in Google Photos™. Follow these steps to do so:

1. In your Document Library, click the *Add* button (+) and select *Google Docs*. The *New Google Docs* screen appears.
2. Click the *Select File* button to open Google's file picker.
3. Use the file picker to select a file from Google Drive™ or Google Photos™.
4. Click *Publish*.

A new file entry appears for the Google document you linked to. You can view the file entry as you would any file entry. The Google document's contents show in the file entry's preview pane. As with any file entry, the *Options* button (⋮) gives you access to the Download, Edit, Move, Permissions, Move to Recycle Bin, and Checkin/Checkout/Cancel Checkout options.

43.10 Metadata Sets

You can define metadata fields that users fill out when they create or edit Document Library files. You do this by creating *metadata sets* and then associating them with document types, which wrap Document Library files and thus apply your metadata fields to the files. Although you apply metadata sets via document types, metadata sets exist independently and you can apply them to any number of document types.

Managing Metadata Sets

To see the available metadata sets, open the *Menu* (▣), expand your site's menu, and navigate to *Content* → *Documents and Media*. Then click the *Metadata Sets* tab. The sets appear in a table.

To select a metadata set, select the checkbox to its left. To select all the sets, select the checkbox in the Management Bar. With one or more sets selected, an X icon appears in the Management Bar. Clicking it deletes the selected metadata set(s). Note that metadata sets don't support the Recycle Bin. If you delete a metadata set, it's gone forever.

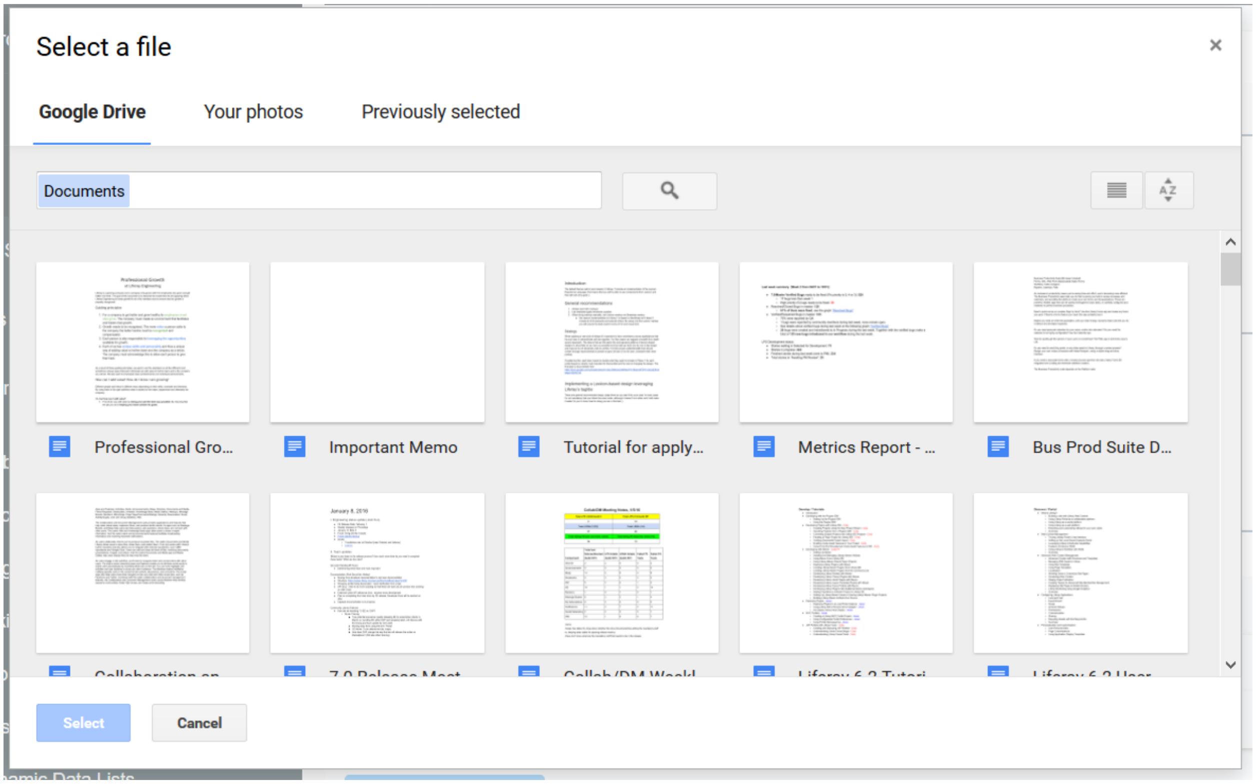

Figure 43.19: You can select files from Google Drive™ or your photos.

The Management Bar also contains other options for managing the metadata sets. The selector menu to the right of the checkbox filters the sets the table displays (it's set to *All* by default). The *Order by* selector orders the sets by Modified Date or ID. The up and down arrows sort the sets in ascending or descending order, respectively. You can also use the Search bar to search for a set.

In the table, each metadata set has an Actions button (⋮) for performing the following actions on that set:

- **Edit**: Edit the set. Alternatively, click the set's name in the table.
- **Permissions**: Configure the set's permissions.
- **Copy**: Copy the metadata set.
- **Delete**: Delete the set.

Creating Metadata Sets

Follow these steps to create a metadata set:

1. From the *Menu* (), click your Site's name and navigate to *Content* → *Documents and Media*. Then click the *Metadata Sets* tab.
2. Click the *Add* button (). The New Metadata Set form appears.
3. Give your metadata set a name.

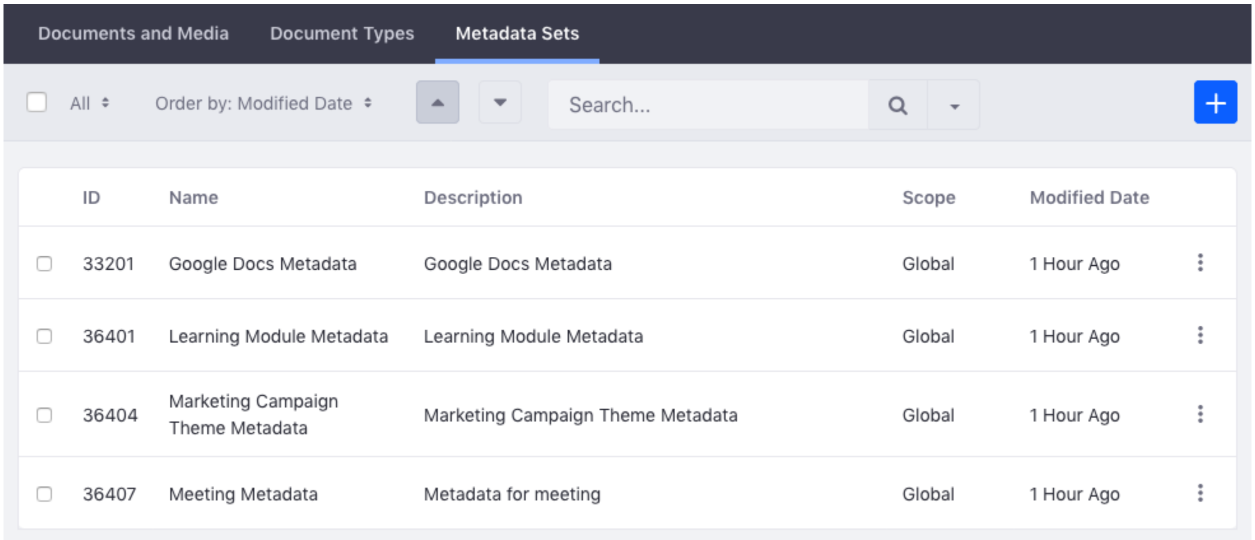

Figure 43.20: The Metadata Sets management window lets you view existing sets and create new ones for applying to document types.

4. Open the *Details* section of the form to give your metadata set a description or select a metadata set to extend (both are optional). To select a metadata set to extend, click the *Select* button for *Parent Metadata Set* and then select the metadata set. When a user creates a document of a document type that uses an extended metadata set, the parent metadata set's fields appear above the extended metadata set's.

5. Add the metadata fields that should be part of this metadata set. To do this, first select the editor's *View* tab and select the *Fields* tab within it. Icons representing the field types are listed on one side and the metadata set's canvas is on the other side. To add a field type to the metadata set, select its icon, drag, and drop it onto the canvas. The field shows on the canvas as for document type users. By dragging a field onto a field that's already on the canvas, you can nest the new field in the existing field. When you mouse over a field on the canvas, the field action icons () appear. Clicking the + icon creates a duplicate of the current field and adds it below the current field. Clicking the trash can deletes the field.

   The following metadata fields are available:

   - **Boolean:** A check box.
   - **Color:** Specifies a color.
   - **Date:** Enter a date. A valid date format is required for the date field, but you don't have to enter a date manually. When you select the date field a mini-calendar pops up which you can use to select a date.
   - **Decimal:** Enter a decimal number. The value is persisted as a `double`.
   - **Documents and Media:** Select a file from a Documents and Media library.
   - **Geolocation:** Specify a location to associate with the document.
   - **HTML:** An area that uses a WYSIWYG editor to enhance the content.
   - **Integer:** Enter an integer. The value is persisted as an `int`.
   - **Link to Page:** Link to another page in the same site.

- **Number:** Enter a decimal number or an integer. The value is persisted either as a `double` or an `int`, depending on the input's type.
- **Radio:** Displays several clickable options. The default number of options is three but this is customizable. Only one option can be selected at a time.
- **Select:** This is just like the radio field except that the options are hidden and must be accessed from a drop-down menu.
- **Text:** Enter a single line of text.
- **Text Box:** This is just like the text field except you can enter multiple lines of text or separate paragraphs.
- **Web Content:** Select web content.

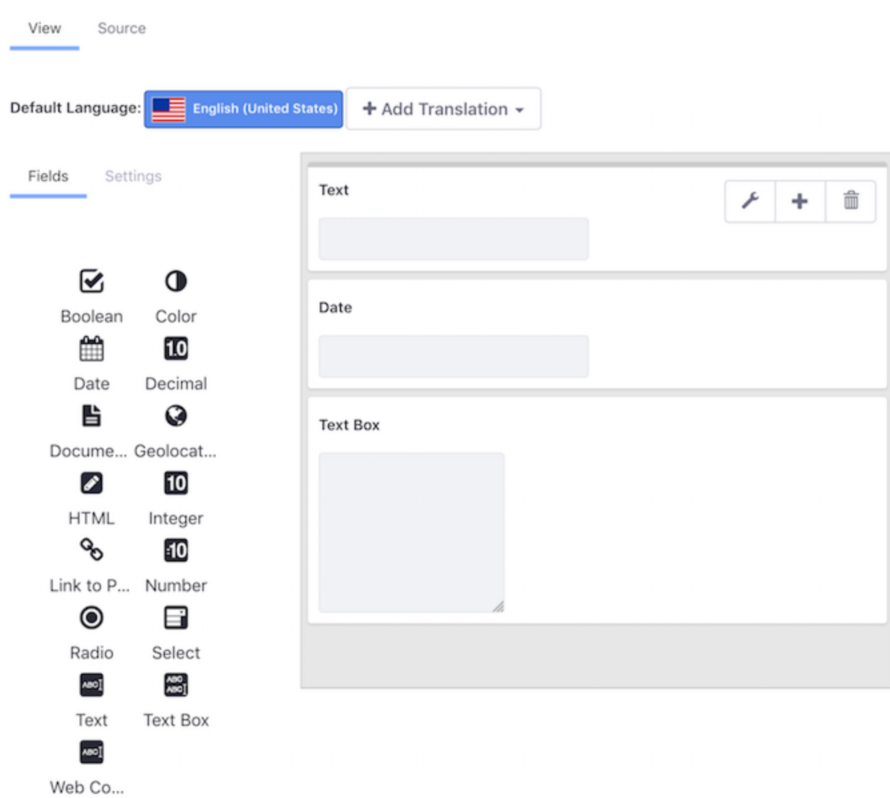

Figure 43.21: Add your metadata set's fields to the canvas.

6. Edit your fields to reflect their intended metadata. For example, a text field's default label is *Text*. If you want to use the text field as a title, for instance, then you should change the field's label to *Title*. To do this, first select the field on the canvas. This automatically selects the *Settings* tab on the left. Alternatively, you can access the Settings tab by clicking the field's wrench icon. To edit a setting value, double-click it in the Settings table and enter the new value.

   Labels, default values, variable names, mouse-over tips, widths, and other settings can be configured for most fields. Some fields have a *Required* setting for specifying whether users must populate the field. If a field's *Repeatable* setting is set to *Yes*, users can add multiple consecutive instances of the field to the document's metadata.

Also note that you can translate each of a metadata set's field values to any supported locales. To specify a field value for a translation, select the flag that represents the locale and enter the field value for the locale.

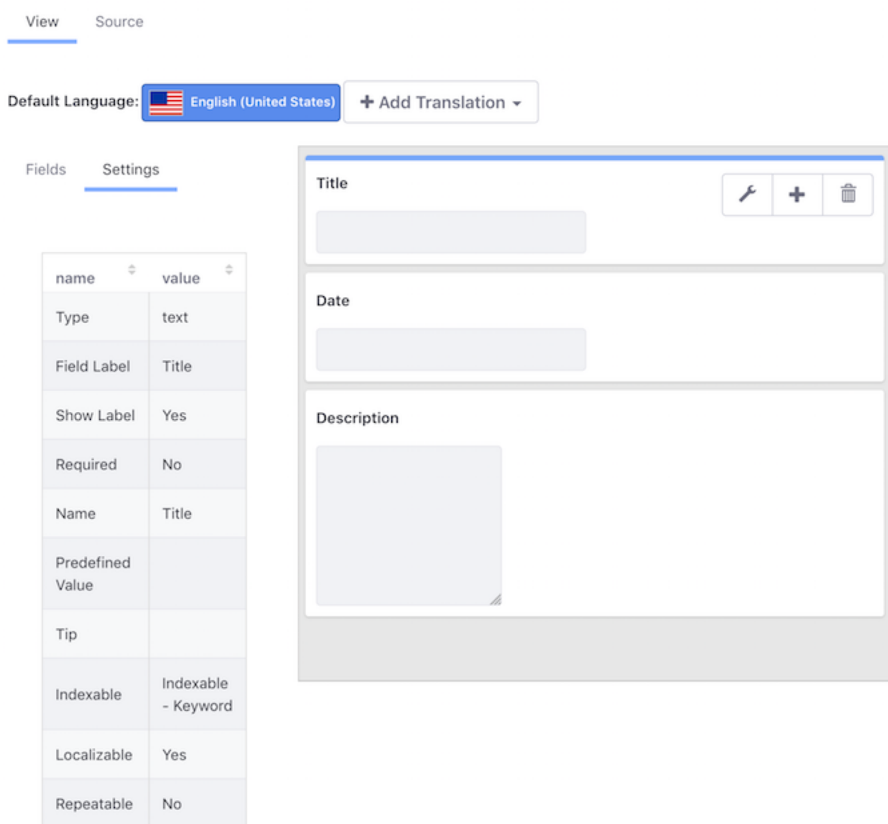

Figure 43.22: Edit your metadata set's fields to match the metadata that you want each field to hold.

7. Click *Save* when you're done specifying your new metadata set.

43.11 Document Types

Document types are made of metadata fields and help users define the purpose of Document Library files. For example, the built-in *Contract* document type contains metadata fields for the effective date, expiration date, contract type, legal reviewer, and more. When users create Document Library files of the Contract document type, they can then populate those metadata fields. Document types also help you integrate files with other features like search and workflow. Search works on file metadata so users can find files faster. You can also apply workflows to specific document types. And you can more cleanly organize document libraries by designating folders to hold particular document types exclusively.

Managing Document Types

To see the available document types, open the *Menu* (), expand your site's menu, and navigate to *Content* → *Documents and Media*. Then click the *Document Types* tab. A searchable table lists the document types. The following actions are available for each document type via its Actions button ( ):

- **Edit**: Edit the document type.
- **Permissions**: Set the document type's permissions.
- **Delete**: Delete the document type. Note that document types don't support the Recycle Bin. Once you delete a document type, it's gone forever.

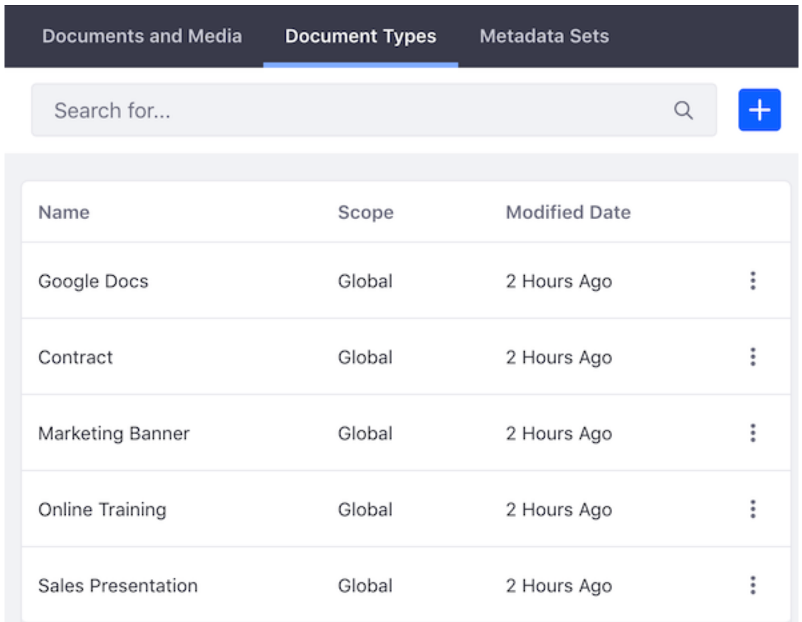

Figure 43.23: The Document Types management window lets you view existing document types and create new ones.

Creating Document Types

Follow these steps to create a document type:

1. From the *Menu* ( ), expand your site's menu and navigate to *Content* → *Documents and Media*. Then click the *Document Types* tab.

2. Click the *Add* button ( ). The *New Document Type* form appears.

3. Give your document type a name and a description.

4. Define the metadata to use with your document type. You do this via these sections in the form:

   - **Main Metadata Fields:** These are tied directly to the document type. They can be created only via the form and can't be used with other document types. You create and edit these metadata fields in the form the same way that you do when creating metadata sets.
   - **Additional Metadata Fields:** Select a metadata set to associate with the document type. Each document type must be associated with one or more metadata set. To differentiate document types that use the same metadata sets, define different main metadata fields.

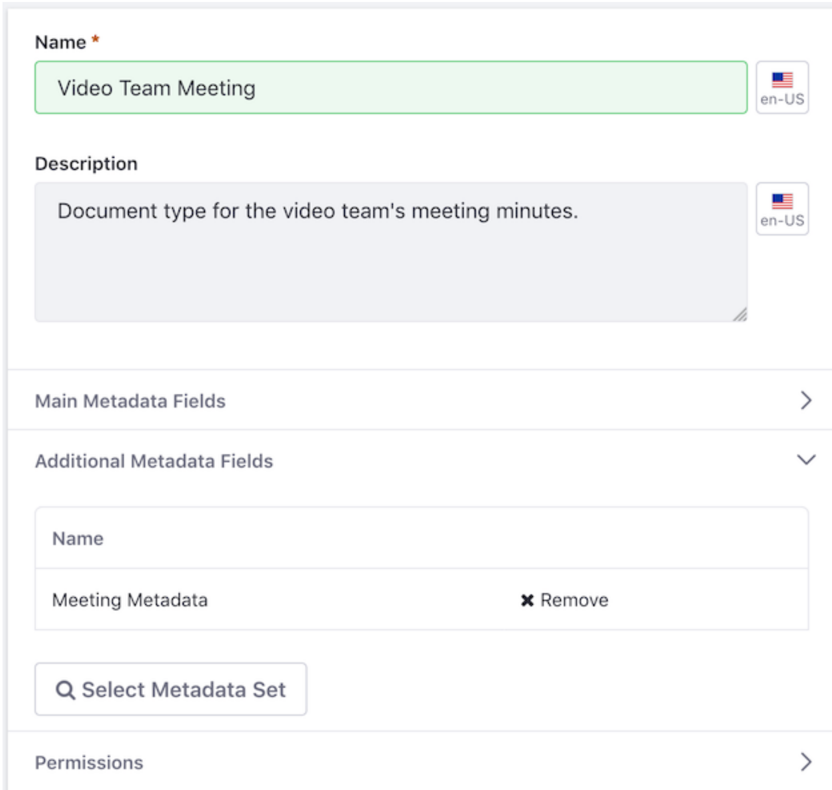

Figure 43.24: Create your new document type.

5. Define your document type's permissions via the form's *Permissions* section. By default, anyone can view the document type, including site guests. You can restrict its view, update, delete, and permissions configuration to site members or the document type's owner.

6. Click *Save* when you're finished specifying your new document type.

Your document type is now available when adding a document via the Documents and Media's *Add* menu. When users create new files of the document type, they're presented with metadata fields to describe the document.

43.12 Store Types

You can change the file system (the *store*) that the Documents and Media library uses to store files. This is configured in the `portal-ext.properties` file by setting the `dl.store.impl=`property. Configuring stores is covered in the Document Repository Configuration guide. Here, you'll consider the ramifications of different stores:

- **Simple File System Store:** Uses the file system (local or a mounted share) to store files. This is the default store.

- **Advanced File System Store:** Nests files into directories by version, for faster performance and to store more files.

- **DBStore (Database Storage)**: Stores files in the Liferay DXP database. The file (stored as a blob) size limit is 1 GB. Use the Simple File System Store or Advanced File System Store to store larger files.

- **S3Store (Amazon Simple Storage)**: Uses Amazon's cloud-based storage solution.

If you must move your files from one store to another, use the migration utility in *Control Panel → Configuration → Server Administration → Data Migration*.

Simple File System Store

The Simple File System Store is the default store. It stores Documents and Media files on the server's file system (local or mounted). This store is heavily bound to Liferay DXP's database. The store's default root folder is `[Liferay Home]/data/document_library`. You can change this via the `dl.store.file.system.root.dir=` property in a `portal-ext.properties` file, or in the Control Panel. For instructions on this, see the Document Repository Configuration guide.

The Simple File System Store uses a local folder to store files. You can use the file system for your clustered configuration, but you must make sure the folder you point the store at can handle things like concurrent requests and file locking. You must therefore use a Storage Area Network or a clustered file system.

The Simple File System Store creates a folder structure based on primary keys in Liferay DXP's database. If, for example, you upload a presentation with the file name `workflow.odp` to a folder named `stuff`, the store creates a folder structure like this:

```
/companyId/folderId/numericFileEntryName/versionNumber
```

- `companyId`: The site's company ID.
- `folderId`: The ID of the Documents and Media folder containing the document.
- `numericFileEntryName`: The document's numeric file entry name.
- `versionNumber`: The document's version number.

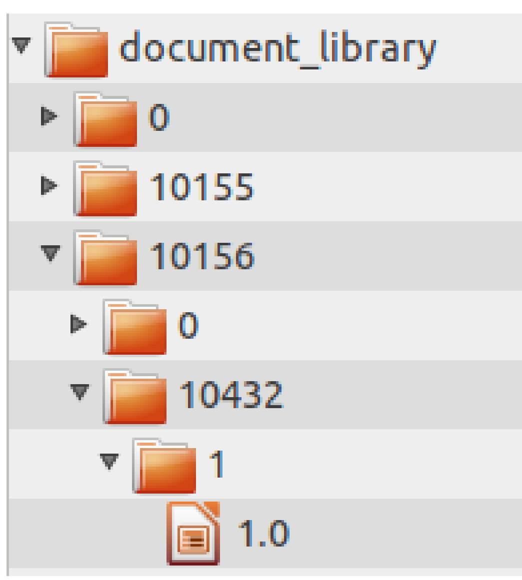

Figure 43.25: The Simple File System Store creates a folder structure based on primary keys in Liferay DXP's database.

Note: Be careful not to confuse a document's numeric file entry name from its document ID. Each has an independent counter. The numeric file entry name is used in the folder path for storing the document, but the document ID is not. The numeric file entry name can be found in the `name` column of the `DLFileEntry` table in Liferay DXP's database; the document ID can be found in the `fileEntryId` column of the same table.

Using the Advanced File System Store

The Advanced File System Store, like the Simple File System Store, saves files to the local file system. It uses a slightly different folder structure, however, and can overcome operating system limitations on the number of files stored in a particular folder by programmatically creating a structure that can expand to millions of files. It alphabetically nests the files in folders. This also improves performance, as there are fewer files stored per folder.

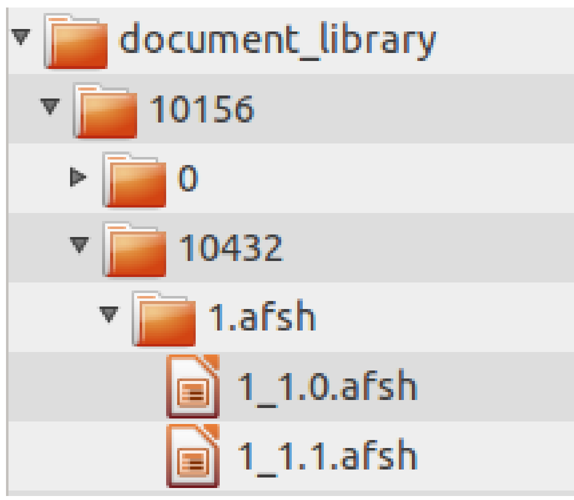

Figure 43.26: The Advanced File System Store creates a more nested folder structure than the Simple File System Store.

The same rules apply to the Advanced File System Store as apply to the Simple File System Store. To cluster it, you must point the store to a network mounted file system that all the nodes can access. That networked file system must also support concurrent requests and file locking. Otherwise, you may experience data corruption issues if two users attempt to write to the same file at the same time from two different nodes.

See the Document Repository Configuration guide for instructions on using the Advanced File System Store.

Using Amazon Simple Storage Service

Amazon's Simple Storage Service (S3) is a cloud-based storage solution that you can use with Liferay DXP. It lets you store your documents to the cloud seamlessly from all nodes.

When you sign up for the service, Amazon assigns you unique keys that link you to your account. In Amazon's interface, you can create *buckets* of data optimized by region. Once you create these to your specifications, follow these instructions to connect your repository to Liferay DXP.

Consult Amazon's S3 documentation for more information.

Chapter 44

Using a SharePoint Repository

The Liferay Marketplace app *Liferay Connector for SharePoint* lets users access SharePoint 2013 and SharePoint 2016 libraries from Liferay DXP's Documents and Media Library. Once this app is installed, you can add a SharePoint repository type to Liferay DXP's Documents and Media Library to access your SharePoint files.

Note: Liferay Connector for SharePoint uses Azure ACS with OAuth 2 for SharePoint server authorization. You must therefore enable HTTPS support in your app server. Consult your app server's documentation for instructions. For example, the required steps for Tomcat can be found in its documentation.

Liferay Connector for SharePoint provides these key benefits:

- Reading/writing documents and folders
- Document check-in, check-out, and undo check-out
- Downloading documents
- Moving folders and documents within the repository
- Getting revision history
- Reverting to a revision

The app uses SharePoint's API, which has these limitations:

- Version history is lost when moving or renaming a file without first checking it out.
- You can't change file extensions; you can only change file names.
- A file's current name propagates to all previous versions.
- The user who checks out a file is the only one who can see the version number of that file's working copy.
- Queries for suffixes or intermediate wildcards convert to queries for containment.
- Comments, ratings, and using a SharePoint folder as a Documents and Media root folder are unsupported.

To use a SharePoint repository in Documents and Media, you must first create an application in SharePoint and authorize it to access the repository. These tutorials walk you through this process.

44.1 Creating a SharePoint Application

The steps covered in this section apply to SharePoint Online. If you're using SharePoint On-Premises, please consult your administrator for the requirements.

Follow these steps to register Liferay DXP as an application in a SharePoint instance:

1. Go to your SharePoint installation's URL: `https://[your-site-name].sharepoint.com/_layouts/15/appregnew.asp`

2. Provide the information below for your app:

   - **Title:** The name displayed in Documents and Media.
   - **Domain Name:** The application's domain name along with the port (e.g., `localhost:8228`)
   - **Redirect URL:** The application's URL. The URL must use HTTPS.

3. Click the two *Generate* buttons to generate a client ID and client secret for Liferay DXP.

   Here's an example configuration:

   - **Client ID:** `1234a56b-7890-1234-5ccc-67d8ea9b0c1c`
   - **Client Secret:** `1ABCDEfGh2IJKLmNoP3QrStuvwX41YzAB+CDEFg20G3=`
   - **Title:** `My Application's Title`
   - **App Domain:** `localhost:8228`
   - **Redirect URL:** `https://localhost:8228/c/document_library/sharepoint/oauth2`

4. Now you must grant Liferay DXP write and search permissions over the SharePoint instance. Other permissions are ignored. Go to `https://[your-site-name].sharepoint.com/_layouts/15/appinv.aspx`.

   In the *APP ID* field, enter the Client ID of the application you just created and click *Search*. Consult Microsoft's documentation for details on how to configure the Permission Request XML. Here's an example configuration that grants the application write and search permissions over the SharePoint instance:

```
<AppPermissionRequests>
    <AppPermissionRequest
        Scope="http://sharepoint/content/sitecollection/web/list"
        Right="Write"
    />
    <AppPermissionRequest
        Scope="http://sharepoint/search"
        Right="QueryAsUserIgnoreAppPrincipal"
    />
</AppPermissionRequests>
```

5. Once you provide the permissions XML, click *Create*.

6. Next, go to *Settings → Site App Permissions* in your SharePoint installation. You can also access this page directly with a URL similar to this: `https://[your-site-name].sharepoint.com/_layouts/15/appprincipals.aspx?Scope=Web`. Find your application in the registration list and note its *app identifier*. You'll use this to create a repository configuration. Here's an example app identifier:

```
i:0i.t|ms.sp.ext|6123d38d-2998-4972-9aaa-71a4da9f3a5a@b9c24ab3-ad34-4943-ab57-729d8c329053
```

Great! Now Liferay DXP is registered as an application in your SharePoint instance. Next, you must create a new SharePoint repository type.

44.2 Creating a New SharePoint Repository Configuration

To connect to a remote SharePoint server you must create a repository configuration:

1. Inside your Liferay DXP instance, open the Control Panel and go to *Configuration* → *System Settings* → *Collaboration* → *Sharepoint OAuth2*.

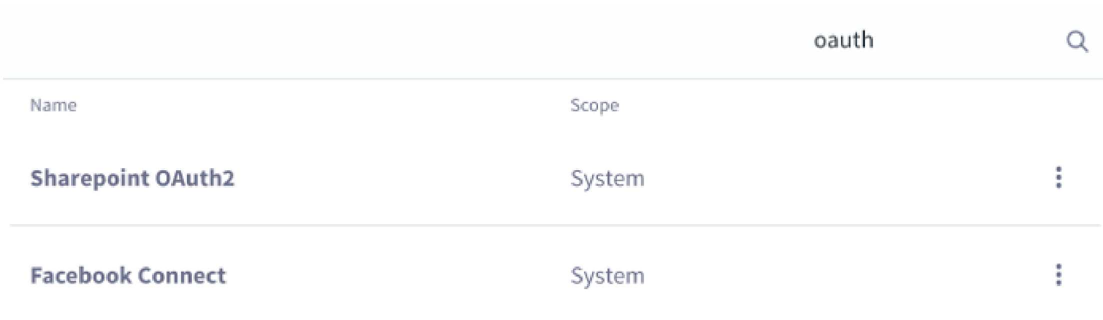

Figure 44.1: Use the *Sharepoint OAuth2* system setting to create a new SharePoint repository configuration.

2. Click the *Add* icon (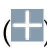) to begin creating a new configuration.

3. In the new Repository form, specify values for the following fields (your SharePoint administrator can provide you with this information):

   Name: The configuration's name.

   Authorization grant endpoint: The URL used to request OAuth2 authorization grants. The URL follows this pattern in SharePoint Online: `https://[your-site-name]/sharepoint.com/_layouts/oauthauthorize.aspx`.

   Authorization token endpoint: The ACS URL. In SharePoint Online with Azure ACS, the URL follows this pattern: `https://accounts.accesscontrol.windows.net/[App ID]/tokens/OAuth/2`.

   Client ID: The client ID.

   Client Secret: The client secret.

   Scope: The permission set required for your tokens. Valid scopes are configured during your app's SharePoint registration.

   Tenant ID: The Tenant ID. If you're using SharePoint Online, you can use the same App ID you embedded in the *Authorization token endpoint*.

   Site domain: The site domain registered for your application.

   Resource: This value depends on the ACS service you use. In SharePoint Online with Azure ACS, the value follows this pattern: `00000003-0000-0ff1-ce00-000000000000/[your-site-name].sharepoint.com@[tenant ID]`.

4. Click *Save*.

Awesome! Now that your SharePoint repository is configured, you can mount it in Liferay DXP's Documents and Media Library.

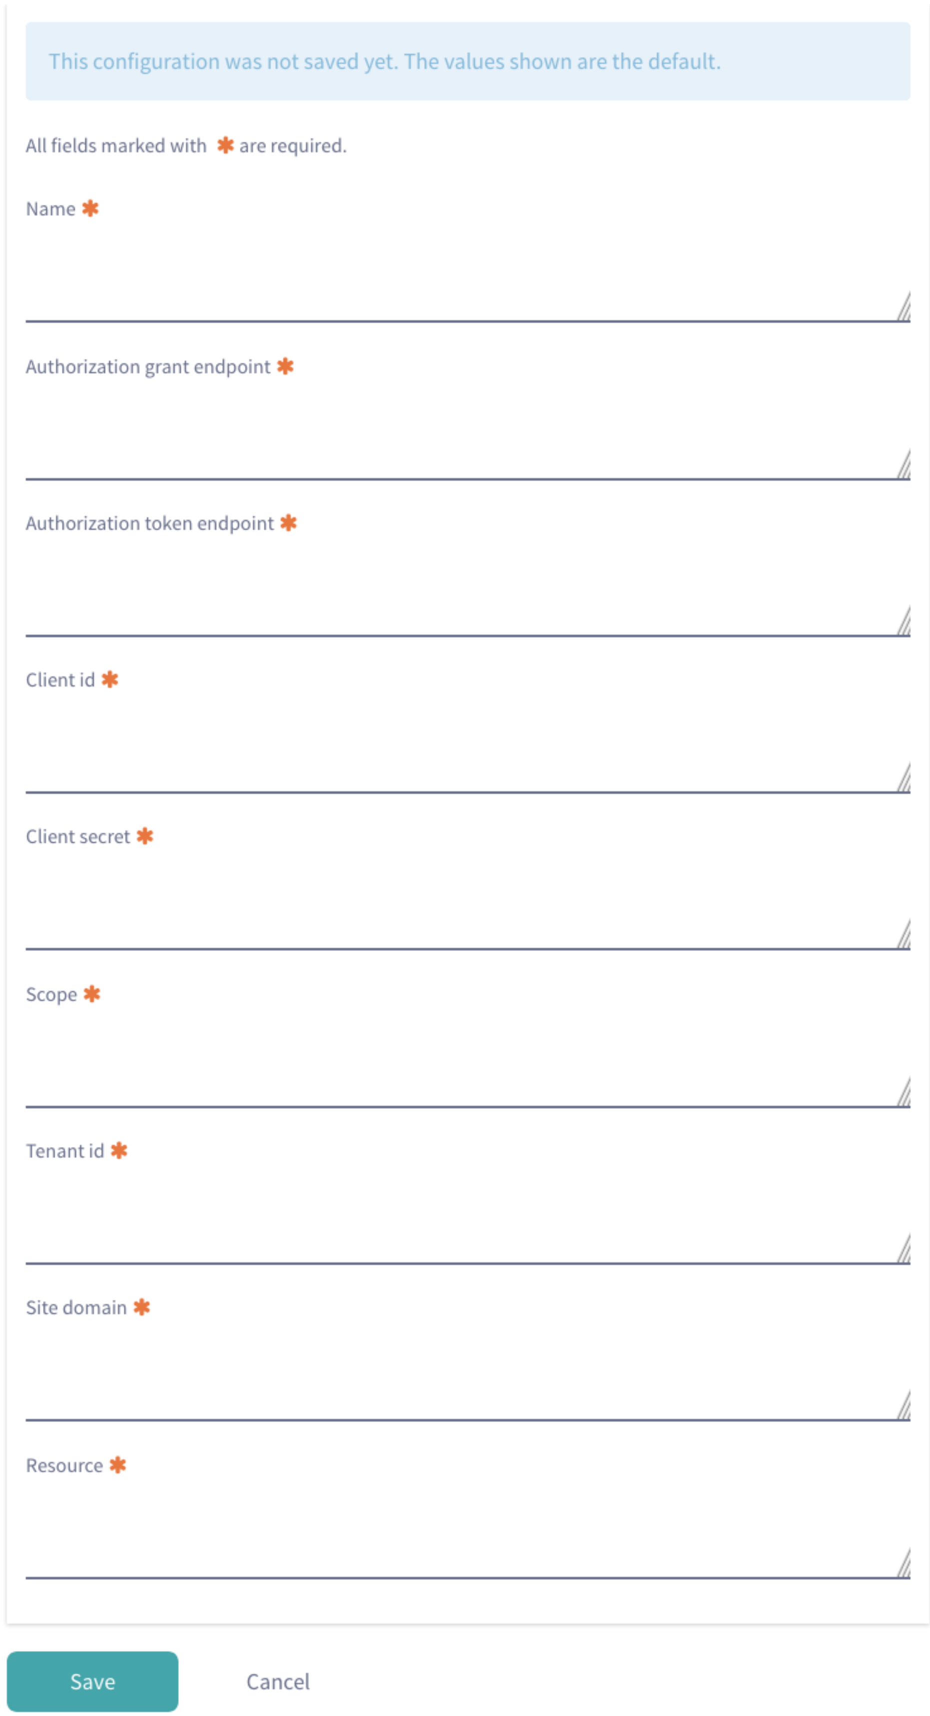

Figure 44.2: The New Repository form is where you specify access to the remote SharePoint server.

44.3 Adding a SharePoint Documents and Media Repository

Follow these steps to add your SharePoint Library to Documents and Media in Liferay DXP:

1. Add the *Documents and Media* application to a page, if you haven't already.
2. Inside the Documents and Media application, click the *Add* icon ( ) and select *Repository*. The New Repository screen appears.
3. Select the repository type for the SharePoint OAuth2 configuration you created. For example, if your configuration is named `Foo`, the repository type is listed as `SharePoint (Foo)`.
4. Specify values for these fields:

   Site Absolute URL: Resolves relative URLs. For SharePoint Online, the value follows this pattern: `https://[your-site-name].sharepoint.com`.

   Library Path: A relative path from the *Site Absolute URL* that points to the SharePoint Document Library you want to mount in Documents and Media (an example path could be `Shared Documents`).
5. Click *Save*.

Liferay DXP's Documents and Media Library now lists your SharePoint repository.

Note: The first time you access a mounted SharePoint repository, you must provide login credentials and also grant permission for Liferay DXP to access the remote SharePoint repository.

Sweet! Now that you've added a SharePoint Repository to Liferay DXP's Documents and Media, you can access and modify SharePoint Library files in your Liferay DXP instance.

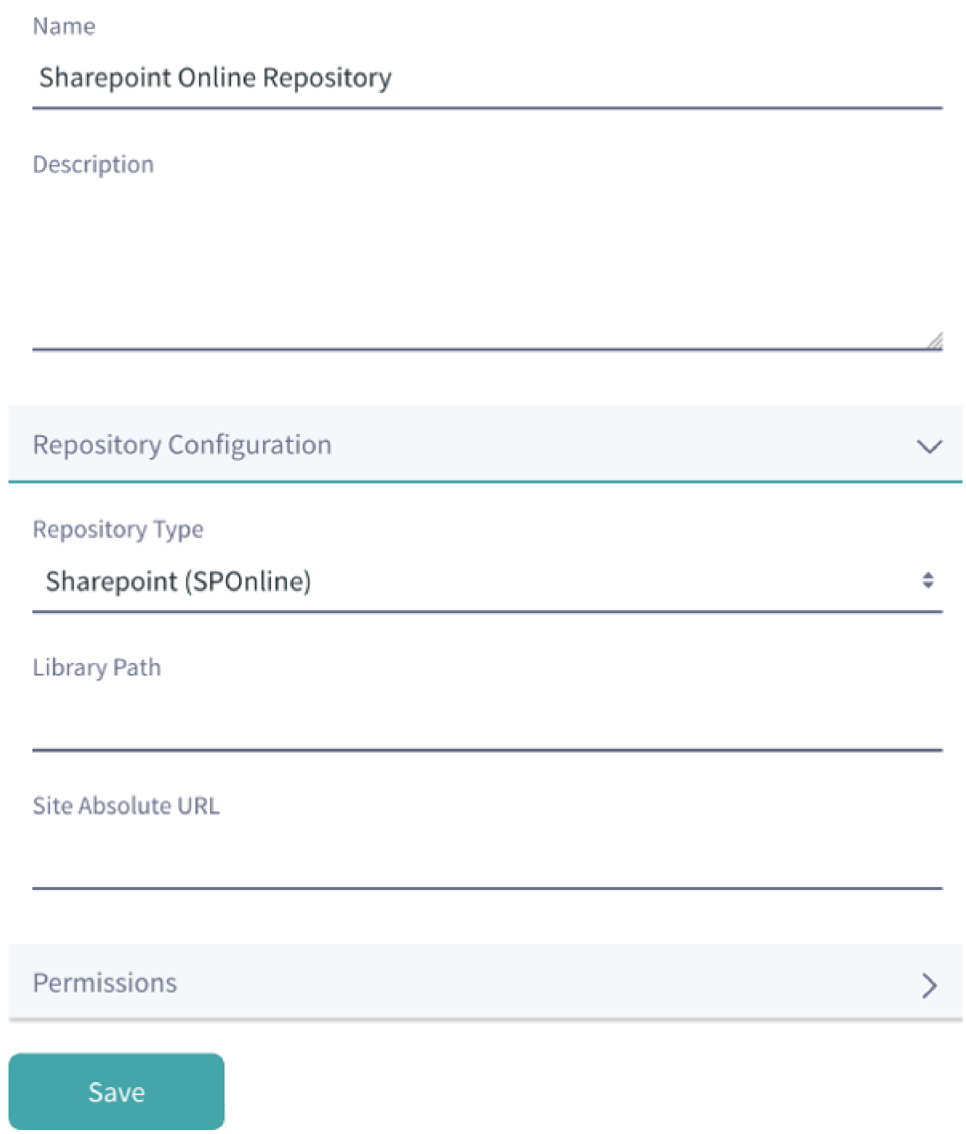

Figure 44.3: The Repository Configuration form is where you specify access to the SharePoint Library you want to use.

Chapter 45

Liferay Sync

Liferay Sync synchronizes files between your Liferay DXP server and users' desktop and mobile environments. With Liferay Sync, users can publish and access shared documents and files from their native environments without using a browser. Clients for Windows and Mac OS desktops and Android and iOS mobile platforms are supported. As users add and collaborate on documents and files, Liferay Sync automatically synchronizes them across all configured Sync clients and your server. Liferay Sync integrates completely with Liferay DXP so that features such as authentication and versioning function in the supported environments. The Liferay Sync desktop client stores files locally so they're always available, even when users are offline. It automatically synchronizes files upon client reconnection. The Liferay Sync mobile client saves storage space on users' devices by downloading only the files they choose.

Chapter 46

Administering Liferay Sync

Before your users can use Liferay Sync with their Sites, you must install and configure it on your Liferay DXP server. The articles here walk you through this and also discuss important topics like preventing accidental file deletion and ensuring Sync security. As such, make sure you thoroughly read each article before letting your users connect to Sync.

Note: To install and configure Liferay Sync on your Liferay DXP server, you must be an administrator.

46.1 Installing Liferay Sync's Prerequisites

Liferay Sync requires that you install the following apps from Liferay Marketplace. Be sure to install them in this order:

1. Liferay Connector to OAuth 1.0a: Enables OAuth 1 in your Liferay DXP installation (note that Sync isn't compatible with OAuth 2). You **must** install this app **before** installing the *Liferay Sync Connector* app. Otherwise, your portal will not be able to start.

2. Liferay Sync Connector: Lets you enable and configure Sync in your Liferay DXP instance. For example, you can disable Sync across the instance or on a site-by-site basis. Note that Sync is enabled by default for all your Liferay DXP instance's sites. You **must** install this app **after** installing the *Liferay Connector to OAuth 1.0a* app. Otherwise, your portal will not be able to start.

Warning: If you install *Liferay Sync Connector* before installing *Liferay Connector to OAuth 1.0a*, your portal will not be able to start.

For instructions on installing Marketplace apps, see the Liferay Marketplace documentation.

Note: The Liferay Sync Security module that Sync requires is included and enabled by default. You can verify this by ensuring that the `SYNC_DEFAULT` and `SYNC_TOKEN` entries are enabled in *Control Panel → Configuration → Service Access Policy*.

If you want to use Sync Connector's default settings and are fine with enabling Sync for all your Sites, you can skip the articles that follow on configuring Sync. However, before directing your users to install and configure the Sync desktop and mobile clients, **make sure to read** this guide's articles on preventing accidental file deletion and ensuring Sync security. You should also **warn your users** about the potential for accidental data loss.

Configuring Sync to Use SSO

If you use an SSO (single sign-on) solution, you must ensure that Sync can access the following URLs without being redirected to your SSO server. Sync can't work without direct access to these URLs. You must therefore whitelist these URLs:

```
http(s)://<portal-address>/c/portal/oauth/*
http(s)://<portal-address>/api/jsonws/sync-web.*
http(s)://<portal-address>/sync-web/*
```

For example, if your installation's address is `https://www.joesblog.com`, then you must whitelist the following URLs:

```
https://www.joesblog.com/c/portal/oauth/*
https://www.joesblog.com/api/jsonws/sync-web.*
https://www.joesblog.com/sync-web/*
```

Sync uses the paths specified in the first URL for communication via OAuth, and the paths specified in the remaining URLs for normal communication with your Liferay DXP installation.

You must also enable OAuth in the Sync Connector app. The next articles cover this, as well as other information on how to configure Sync.

46.2 Configuring Liferay Sync

Sync Connector lets you manage how, or if, clients connect to your Liferay DXP server. You can also configure default file permissions on a per-Site basis, and manage the devices that connect to your Liferay DXP instance. To access Sync Connector, select *Control Panel → Configuration → Sync Connector Admin*.

Sync Connector Admin has three tabs:

1. **Settings:** Control Sync's general behavior. These settings apply globally to Sync.

   Allow the use of Sync?: Whether Sync is enabled.

   Allow users to sync their personal Sites?: Whether users can sync data with their personal Sites.

   Allow LAN Syncing?: Whether desktop clients attempt to download updates from other desktop clients on the same local network before downloading from the server. This can help reduce server load and increase data transfer speeds. Note that LAN syncing only works with clients that also enable it.

   Max Connections: The maximum number of simultaneous connections each client is allowed per account. For example, if Max Connections is three, a client can simultaneously upload or download up to three files for each account. Note, this setting operates on a per client basis.

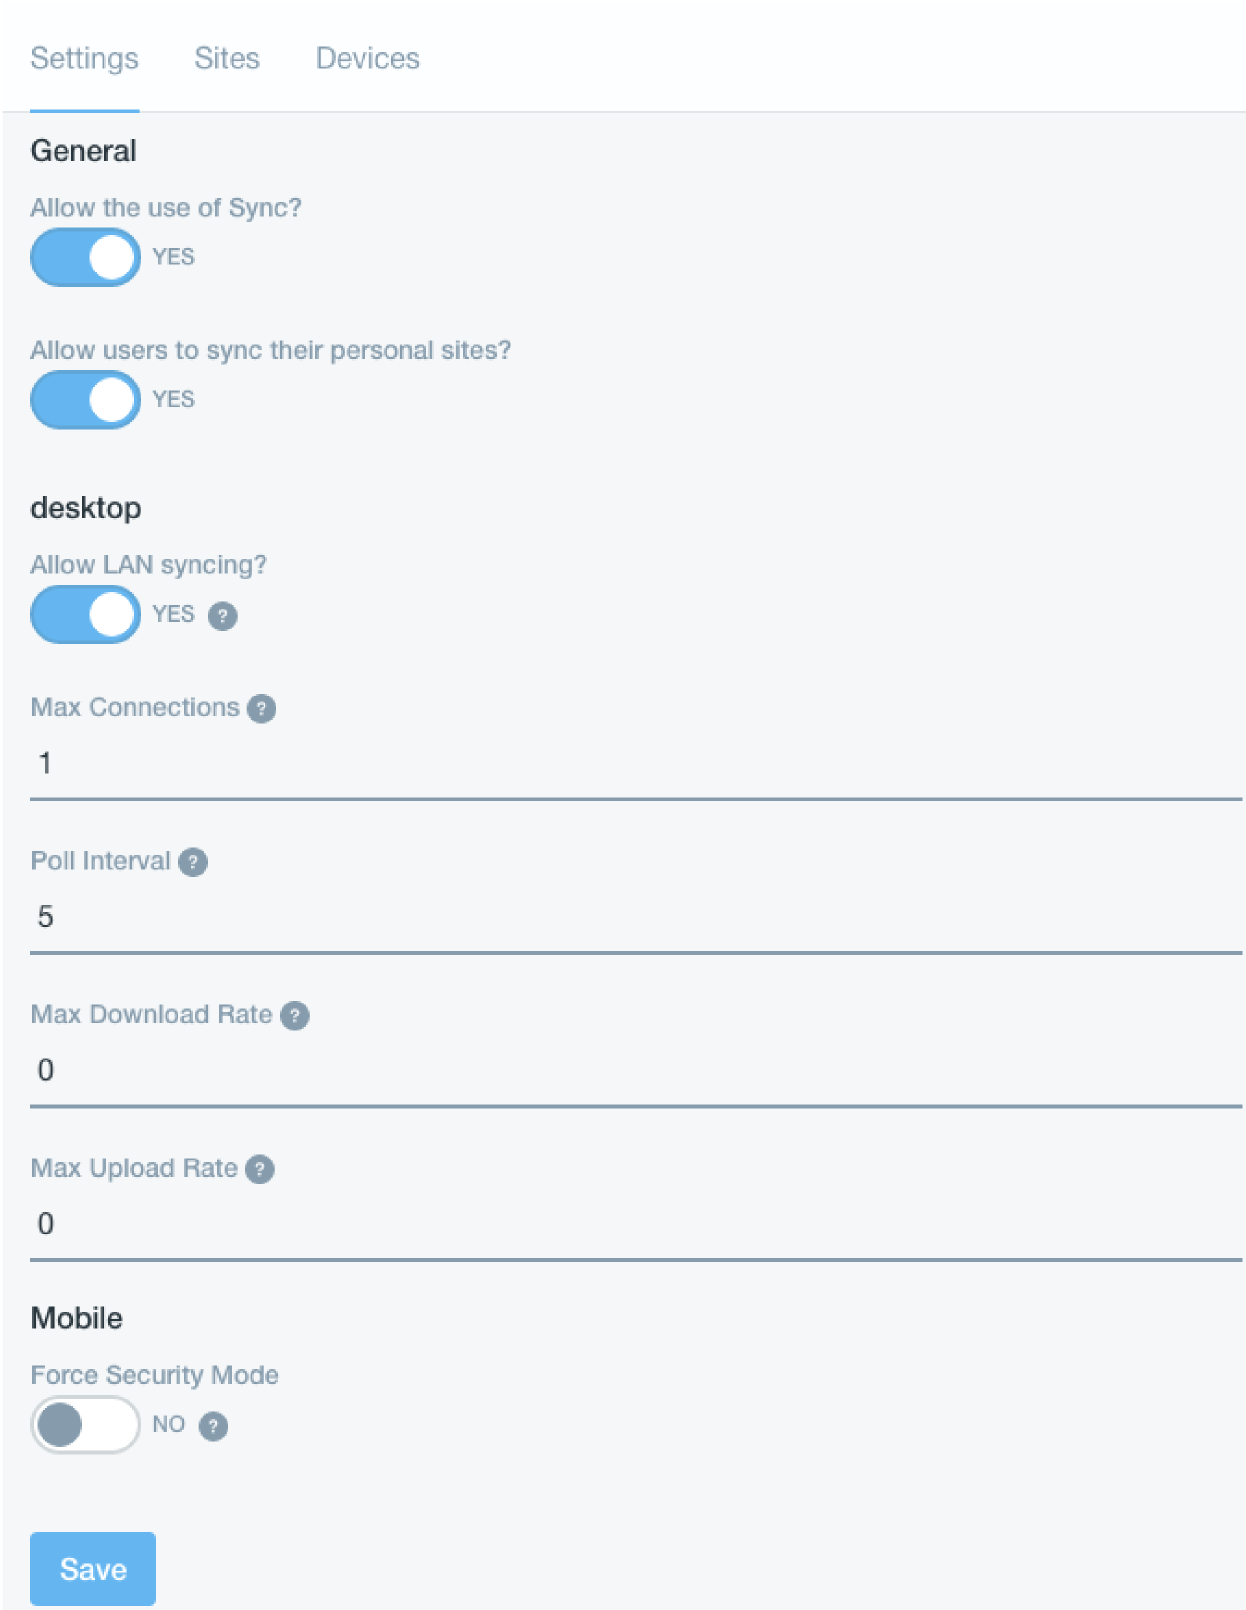

Figure 46.1: The Control Panel's Configuration section contains Sync Connector Admin.

If Max Connections is set to three and a user has two clients connected to an account (which is possible if Sync is installed on two different machines), then the user is effectively allowed six simultaneous connections. While increasing Max Connections can speed up file transfers it also places a heavier load on the server. *Max Connections* is set to one by default.

Poll Interval: The frequency in seconds that clients automatically check the Liferay DXP instance for updates. For example, if set to ten, connected clients check the instance for updates every ten seconds. The default Poll Interval is five.

Max Download Rate: The maximum transfer rate, in bytes, at which clients can download. A value of 0 specifies no limit. This setting takes precedence over clients' download rate setting.

Max Upload Rate: The maximum transfer rate, in bytes, at which clients can upload. A value of 0 specifies no limit. This setting takes precedence over clients' upload rate setting.

Force Security Mode: Whether to force security mode on mobile clients. Security mode encrypts Sync files on the device and requires a passcode when accessing the Sync mobile app.

2. **Sites:** Control Sync on a per-Site basis.

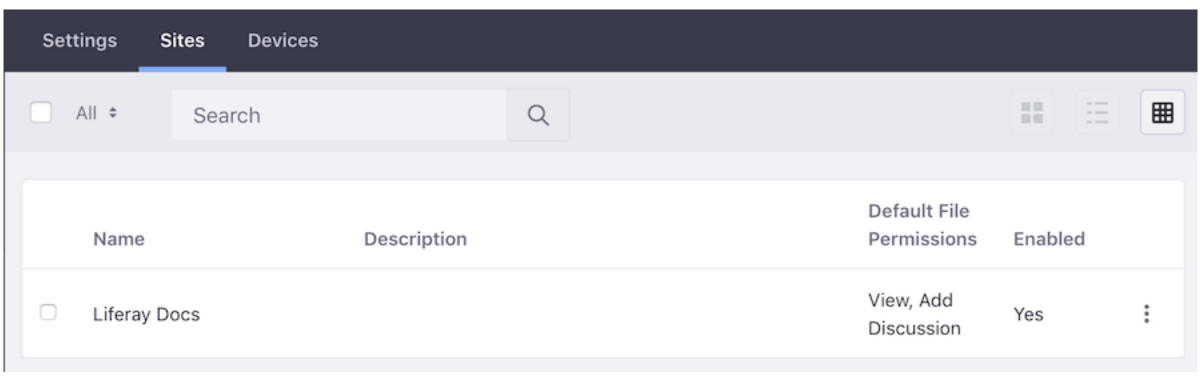

Figure 46.2: Sync Connector Admin's Sites tab lets you manage Sync on a per-Site basis.

For each Site in the Liferay DXP instance, the Sites tab lists each Site's default file permissions (more on this in a moment) and whether Sync is enabled for that Site. Sync is enabled by default for all Sites. To disable Sync for a Site, click the Site's *Actions* button ( ) and select *Disable Sync Site*. **Please use caution** when disabling Sync for a Site, as doing so **deletes** files for that Site from the Sync clients. Disabling Sync for a Site, however, doesn't affect the Site's files on the server.

```
**Warning:** Disabling Sync for specific Sites from Sync Connector Admin can
result in data loss across clients. If Sync is disabled for a Site users are
currently syncing, any files in the clients' sync folders for that Site are
automatically deleted from their clients. If a user is offline when Sync is
disabled for a Site, any offline changes or additions they make are deleted
upon client reconnection.
```

```
You can enable Sync for a Site by selecting *Enable Sync Site* from its
Actions button. Make sure that each Site for which Sync is enabled has a
Documents and Media app on at least one of its pages. If a Site doesn't have
the app on any of its pages and users click the *Open Website* link from
their Sync menus, the error message *The requested resource was not found*
appears.

The Sites tab also sets default file permissions for files uploaded from
Sync clients. The process for setting permissions is nearly the same as for
enabling or disabling Sync for Sites. To set the default file permissions
for a single Site, click its Actions button and select *Default File
Permissions*. This lets you select the default file permissions for that
Site. Click *Choose* for the permissions you want to use.

![ Click *Choose* to select the default file permissions for a Site in Sync.](./images/sync-admin-03.png)

To set the default file permissions for several Sites, select the checkboxes
for the Sites, click the *Default File Permissions* link that appears above
the table, and select the permissions you want to use. Default file
permissions might behave differently than you'd expect. They control *only*
the permissions for new files uploaded through the Sync clients; they don't
affect permissions for uploading or restrict document owners (the user who
originally uploaded a document) in any way. For example, even if you set a
Site's default file permissions to View Only, that Site's users can still
upload new documents to the Site. The file's owner has edit permission; the
rest of the Site's users have the View Only permission.
```

3. **Devices:** View and manage the devices registered with Sync.

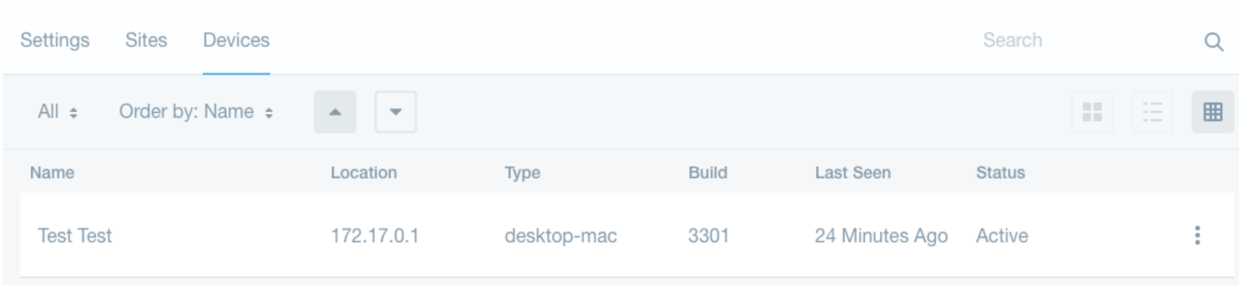

Figure 46.3: Sync Connector Admin's Devices tab lists all the devices Sync has registered.

Each row in the Devices tab's table represents a device. The *Name* column lists the user that registered the device. The remaining columns list each device's location, client type, client build number, last connection date, and status. Each device's Actions button ( ) manages that device. You can change a device's status from Active to Inactive by selecting *Actions* → *Disable Sync Device*. Inactive devices can't sync. Inactive mobile devices also can't access local Sync files. Once a device is Inactive, you can erase Sync files from it by selecting *Actions* → *Wipe Sync Device*. This also signs the device out and removes the account from the client. If the device is offline, this happens when it tries to reconnect. The Actions menu also enables or deletes an Inactive device. Deleting a device only removes it from the list of registered devices; it can still reconnect and reregister.

46.3 Preventing Accidental File Deletion in Liferay Sync

Liferay Sync's power rests in its ability to propagate between Liferay DXP and connected Sync clients. When a user deletes a file from a connected client, Sync also deletes the file in the instance and in any other connected clients. Likewise, if a user deletes a file in the instance, Sync also deletes the file in all connected clients. In other words, anywhere a user deletes a file, Sync deletes it *everywhere*. But don't worry: Liferay DXP's Recycle Bin is enabled by default and lets you recover deleted files. You can access the Recycle Bin from each Site's *Site Administration* menu.

Warning: Liferay Sync automatically propagates file and folder deletion through the Liferay DXP server and in all connected clients. If an instance or Site administrator disables the Recycle Bin, deleted files can't be recovered.

Liferay DXP instance and Site administrators can, of course, disable the Recycle Bin. Disabling the Recycle Bin in a Site, however, leaves the Site vulnerable to accidental file deletions that propagate through Sync.

46.4 Ensuring Liferay Sync Security

As an administrator, you have a stake in the security of all connections to and from your servers. As long as Liferay DXP is configured to use HTTPS, Sync clients use user-supplied credentials to communicate securely. Users can only access the documents and Sites they're permitted to access. To support Security Mode in the Sync mobile client and securely transmit files, your Liferay DXP server must also use SSL. The next section demonstrates how Sync's permissions work with your Liferay DXP instance's permissions.

Liferay Sync Permissions Demonstration

Sync uses Liferay DXP's default permissions to determine files and folders to sync with the user's devices. It can only sync files a user can access. After installing the desktop Sync client, follow the steps below to test this functionality.

First, enter `classified information` into a new text file and save it on your desktop as `secret.txt`. Then use your browser to sign into Liferay DXP and create a new user with the user name *secretagent* and the email address *secretagent@example.com*. Give this user a password and then create a new private Site called *Secret Site*. Create a page on the Site and add the Documents and Media app to it. Then add the secretagent user to the Secret Site and grant the *Site Administrator* Role to the user. Log in as secretagent and navigate to the Secret Site. Then upload the `secret.txt` document to the Documents and Media app. Make sure you also have a user that isn't a member of the Secret Site and therefore doesn't have access to any of its documents through Sync. If you don't have such a user, create one now.

Next, configure your Liferay Sync client to sign in with the secretagent user's credentials and sync with the Secret Site. Open the Liferay Sync menu from the system tray and select *Preferences*. In the *Accounts* tab, click the plus icon at the window's bottom left to add an account. Provide the secretagent user's credentials and uncheck all Sites except the Secret Site. Now confirm that Sync downloaded the `secret.txt` file to your new Sync folder. Open it and check that it contains the text `classified information`. Next, use Sync to connect to your Liferay DXP instance with the user that doesn't belong to the Secret Site. The file doesn't sync because this user isn't a Site member.

Now go to Sync Connector Admin and set the Secret Site's default file permissions to View Only. Create a new user, add it to the Secret Site, and add its account in your Liferay Sync client. As with the secretagent user, Sync downloads the `secret.txt` file to this user's local Sync folder because the user is a member of the Secret Site. Now edit and save this file. Even though you can edit and save it locally, the edits aren't synced because the Site's default file permissions are View Only. After attempting the sync, a red *x* appears next to the file in the local Sync folder. Right click the file to see the error. It confirms the user doesn't have the required permissions.

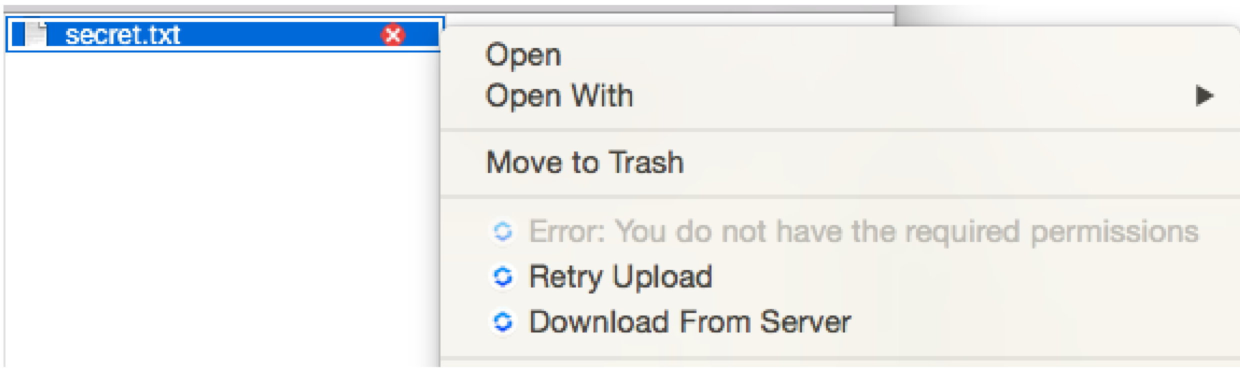

Figure 46.4: The upload error occurs because the user only has permission to view files.

To confirm that the error didn't propagate through Sync, open the file in the secretagent user's local Sync folder. It still contains the original text. Likewise, the original file remains in the Site's Documents and Media portlet. To get rid of the error in the other user's local Sync folder, return there and then right click the file and select *Download From Server*. This replaces the file with the latest file in the Liferay DXP instance.

Now edit `secret.txt` in the secretagent user's local Sync folder. When you check the file in the other user's local Sync folder and in the Liferay DXP instance, notice that Sync propagated the edits. The changes were propagated because the secretagent user owns the file in the instance. Owners can do anything with their files, even when the Site's default file permissions are set to View Only.

CHAPTER 47

USING LIFERAY SYNC ON YOUR DESKTOP

Liferay Sync synchronizes files between your Liferay DXP Sites and desktop devices. It lets you work with your files without using a browser. The Sync clients also ensure that the files are updated with the latest changes made by other users. To use Liferay Sync in your desktop environment, you must install the Sync desktop client. It's currently available for Windows and Mac OS. The Sync client stores files locally so that they're always available, even when you're offline. Files are automatically synchronized upon your client's reconnection to your Liferay DXP server.

On your desktop devices, Liferay Sync creates a new folder structure that it uses to synchronize files. You can treat the files the same as you do any others. Credentials, Sync folder location, and other options are configured in the client. Also, native desktop notification events inform you of what Sync is doing. The native menu and task bar integration keep Sync controls within easy reach.

This guide walks you through setting up and using the Liferay Sync client on your desktop. Before proceeding, check your Liferay DXP instance or Site administrator to ensure that Sync is enabled for your Sites.

47.1 Installing and Configuring the Desktop Liferay Sync Client

You can download the desktop client from the Liferay Sync downloads page. Note that you'll need a Liferay account for this. Once you've downloaded the appropriate desktop client for your operating system, installing Liferay Sync on Windows or Mac OS is straightforward.

Installing the Liferay Sync Desktop Client

To install the Liferay Sync client on Windows, you must have administrator privileges. Upon launching the Windows application installer, you're prompted to choose an install location. Select an appropriate location and click *Install*. Sync automatically starts after the installation finishes. The first time Sync runs, you must configure it to connect and sync with Liferay DXP. The configuration steps are shown below.

Note: You can upgrade previous versions of the desktop Liferay Sync client to version 3.0. When doing so, however, you must set up your account again in the new version of the client. Prior to upgrading, it's typically best to shut down Liferay Sync, backup files from your local Sync folder, and delete that folder.

The Liferay Sync client for Mac is packaged in a DMG file. Double-clicking on the DMG file mounts it as a disk image and opens a window showing the image's contents. To install Sync, drag the Liferay Sync icon to your Applications folder. Once it's installed, run it from the Applications folder. If you're running Mac OS X 10.9 or lower, you're prompted for your machine's administrator credentials to install the Finder icon/context menu tool. This prompt only appears when installing or upgrading the tool.

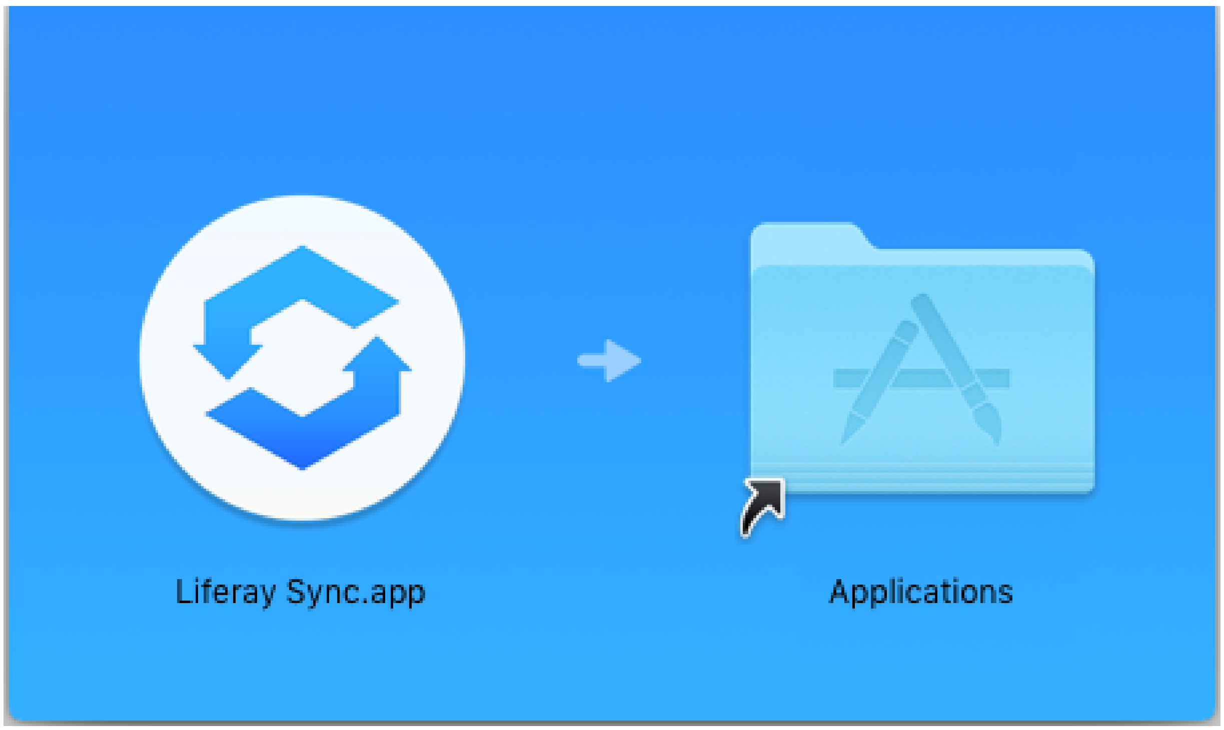

Figure 47.1: Drag the Liferay Sync icon to the Applications folder.

Next, you'll configure the Sync client.

Configuring the Liferay Sync Desktop Client

Now that you've installed Sync, you're ready to configure it! The configuration steps for Sync on Windows and Mac are identical.

1. Open Sync and enter your server's address along with your account credentials. Click *Sign In* when you're finished.

   When connecting to a server via HTTPS, an error appears if the certificate can't be verified. Choosing *Proceed Anyway* bypasses verification and leaves the connection open to compromise. Liferay Sync attempts to read the certificates specified in the Java Control Panel (see section 20.4.5). If Java isn't installed, you can also put your certificates in `[user.home]/.liferay-sync-3/certificates`. Liferay Sync trusts all certificates in this folder.

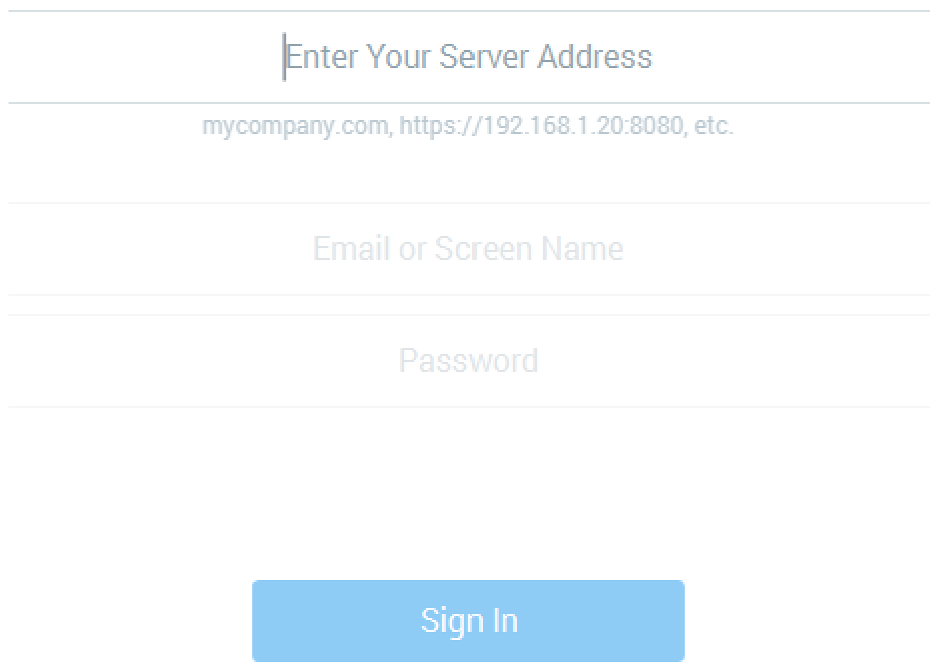

Figure 47.2: The first time you run Liferay Sync, you must tell it how to communicate with your Liferay DXP server.

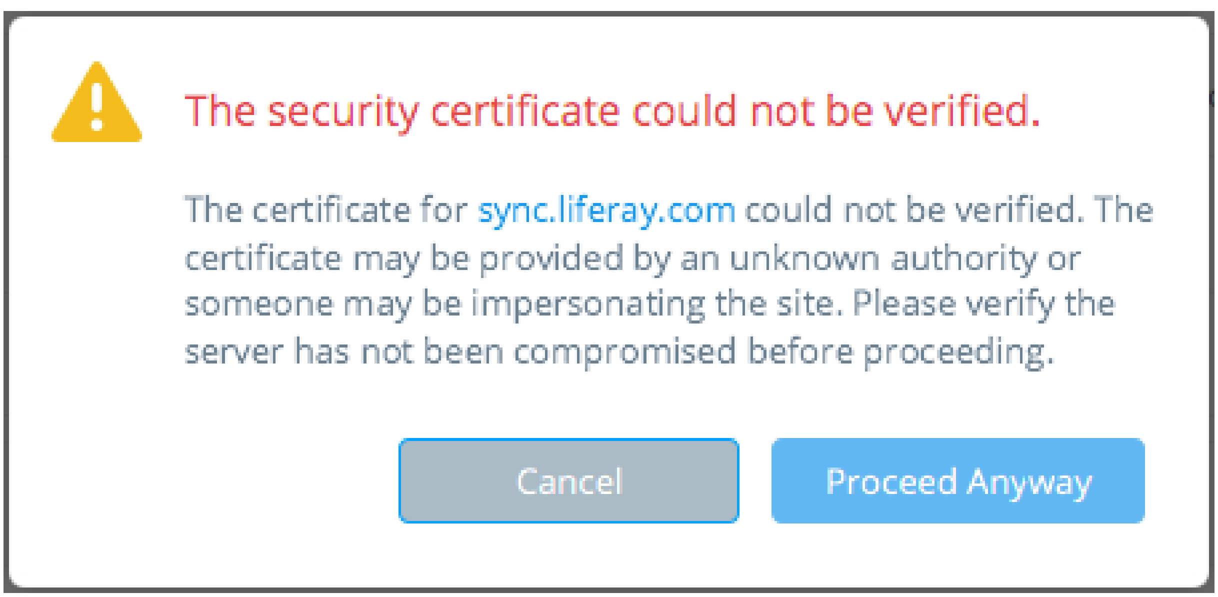

Figure 47.3: When connecting over HTTPS, Liferay Sync produces an error if it can't verify the security certificate. Choosing *Proceed Anyway* bypasses verification and leaves the connection open to compromise.

2. Select the Sites you want to sync with. You can search for a Site in the *Search* bar above the Site list. If you want to sync all the subfolders of your selected Sites, click *Proceed* and move on to the next step.

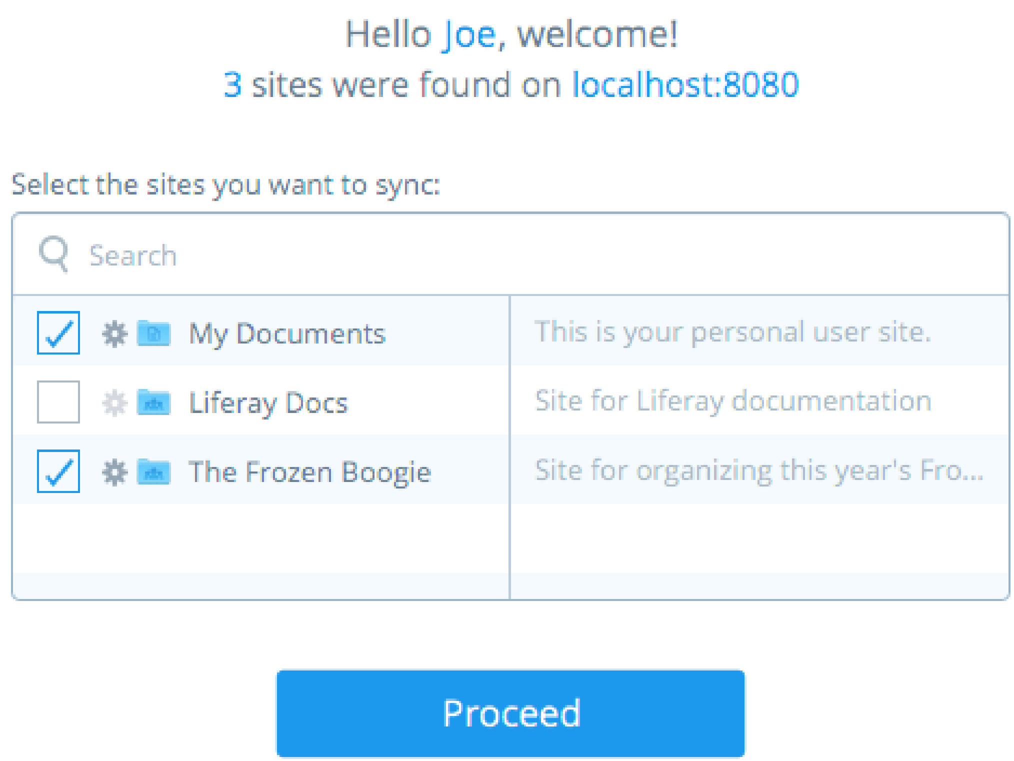

Figure 47.4: Select the Sites you want to sync with. Clicking a Site's gear icon opens another window where you can choose to sync with only specific subfolders in that Site.

To sync only specific folders in a Site, first click the Site's gear icon. In the window that appears, all folders are selected by default. Unselect the folders you don't want to sync with. Unselecting a subfolder causes the parent folder's checkbox to show a minus sign, indicating that you haven't selected all of the parent folder's subfolders. To sync only the documents at the top of a folder's hierarchy (no subfolders), unselect all of that folder's subfolders. You can also do this by clicking the folder's checkbox until the minus sign appears. Click *Select* when you're finished with your selections, and then click *Proceed* to move on to the next step.

3. Specify the local folder to sync. This folder is used exclusively for Sync: Sync creates it and it must not conflict with any existing local folder. The Sync folder's default name is the instance's host name, and its default location is the user's documents folder. For example, since the instance in the following screenshots runs locally at the address `http://localhost:8080/`, Sync creates a Sync folder named *localhost* in the user's documents folder. You can, of course, specify any unique name and location for the Sync folder. Click *Start Syncing* to begin syncing files.

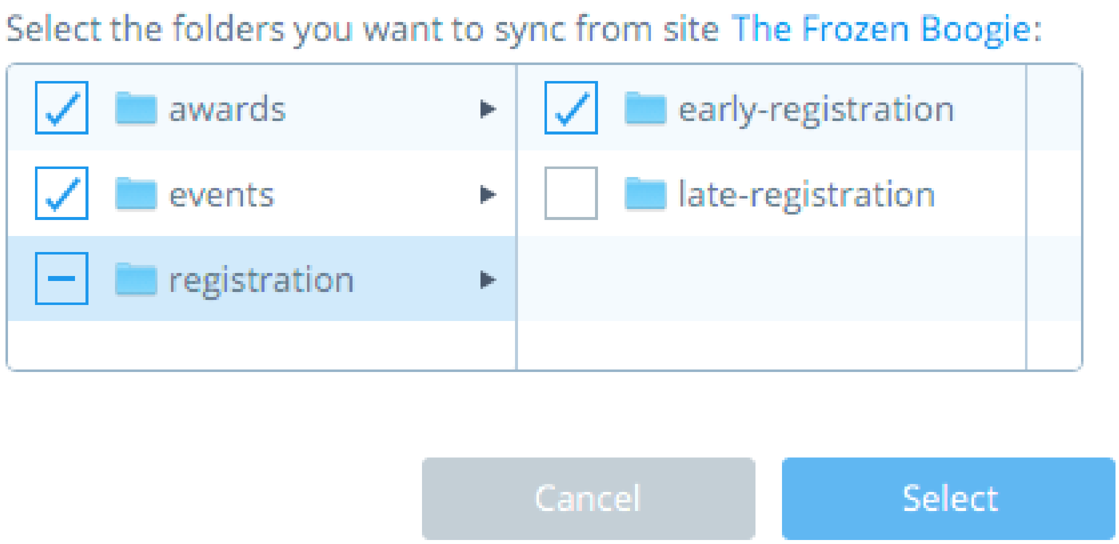

Figure 47.5: Choose the Site's subfolders that you want to sync with. The checkbox with the minus sign indicates that not all of the *registration* folder's subfolders are selected.

```
**Note:** Syncing to network drives is not supported because Liferay Sync
can't reliably detect local file changes on such drives.
```

```
![ Specify your local Sync folder's name and location.](./images/sync-setup-03.png)
```

4. Celebrate! You've successfully set up Liferay Sync! Sync congratulates you on setting it up and begins to sync files from the Sites you selected to your local Sync folder. Note, completing the initial synchronization may take a significant amount of time, depending on the amount of data being transferred. You can safely close the window as syncing continues in the background. To view the local Sync folder, click *Open Folder*. To open Sync's preferences, click the small gray text *advanced setup* near the top-right.

47.2 Using the Liferay Sync Desktop Client

When Liferay Sync is running, its icon appears in your task bar (Windows) or menu bar (Mac). Clicking this icon opens a menu that lets you work with and manage Liferay Sync.

The top of this menu shows your Sync status. If all your selected Sites are synced, then your status is *Synced*.

Below your Sync status, the menu lists three shortcuts for accessing your files:

Open Sync Folder: Select a Site to open its local Sync folder.

View Website: Select a Site to view the page in Liferay DXP that contains its Documents and Media app.

Recent Files: Lists recently created and modified files in the repositories you can access.

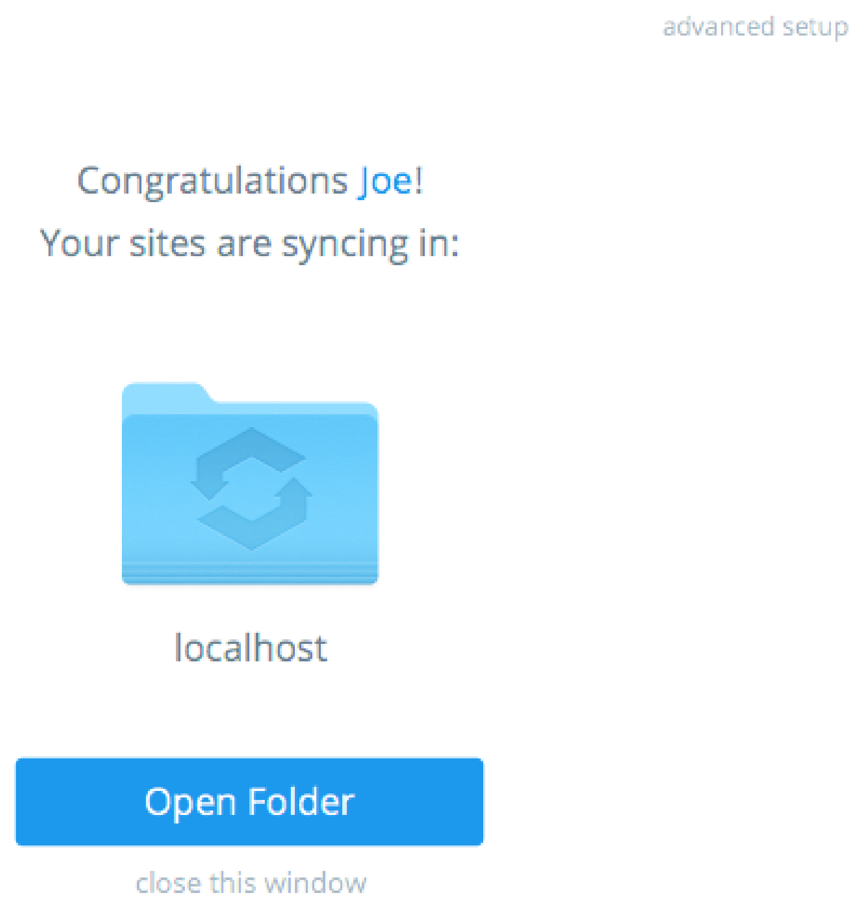

Figure 47.6: Congratulations, you've successfully set up Liferay Sync!

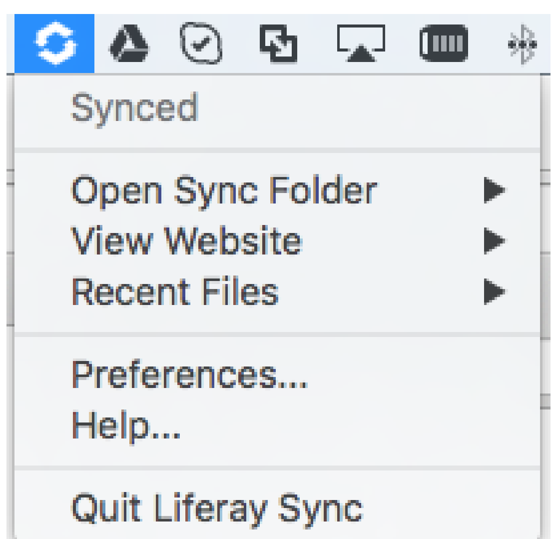

Figure 47.7: The Liferay Sync menu in the Windows task bar and Mac menu bar gives you quick access to Sync.

Note that if you sync with two or more Liferay DXP instances, Sync shows each at the top of the menu instead of your Sync status. Mouse over each instance to reveal a submenu with that instance's Sync status and file shortcuts.

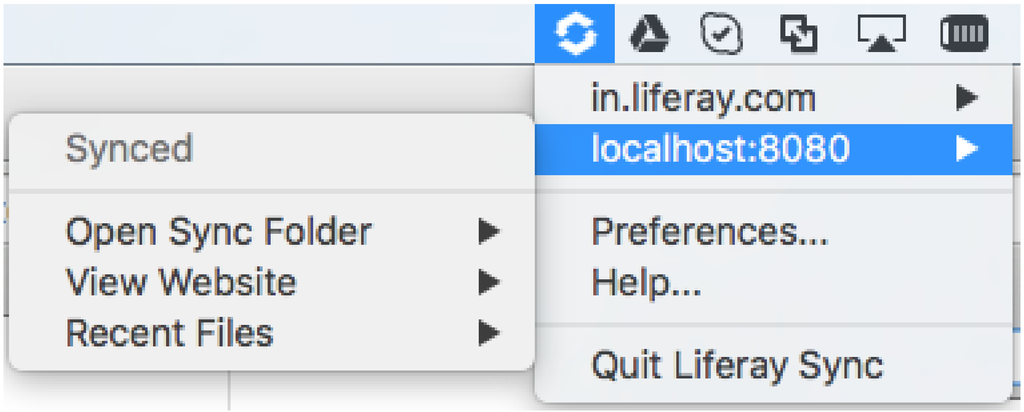

Figure 47.8: When you sync with more than one Liferay DXP instance, Sync shows submenus for each.

Finally, regardless of how many Liferay DXP instances you sync with, the menu lists the following three options:

Preferences: Open Sync's preferences.

Help: Open Sync's documentation.

Quit: Shut down Sync on your machine.

Next, you'll learn how to use Sync's preferences to control how Sync functions on your machine.

Using Sync Preferences

You can use Sync's preferences to add/remove Liferay DXP instances to sync with, edit connection settings, and control Sync's basic behavior. Open Sync's preferences by clicking the Sync icon in the task bar (Windows) or menu bar (Mac OS) and selecting *Preferences*. A preference screen for your instance accounts displays. This is the *Accounts* tab in *Preferences*.

The *Accounts* tab contains the following options:

Accounts: The accounts you sync with. When you select an account, the Sites you have permission to sync with are shown on the right under *Syncing Sites*. You can use the plus, minus, and pencil icons at the bottom of the account list to add, delete, or edit an account, respectively. You should use caution when deleting an account from your Sync client, as doing so also deletes any local files and folders for that account. Adding an account takes you through the same set of steps you used to set up the Sync client. Click here for instructions on this.

Syncing Sites: The Sites you have permission to sync with for the selected account. The Sites you currently sync with are shown under *Selected Sites*. Other Sites available for syncing are shown under *Unselected Sites*. To change the Sites you sync with, click the *Manage Sites* button. The window that appears lets you select and/or unselect Sites to sync with. This window is identical to the one that appeared when you first configured the client. Click here and see step two for instructions on using it. Use caution when de-selecting Sites. De-selecting a Site deletes its folder on your machine.

Location: The selected account's local Sync folder location. Click the *Change* button to change this folder's location.

The Preferences menu's *General* tab contains settings for the Sync client's general behavior. It lists the following options:

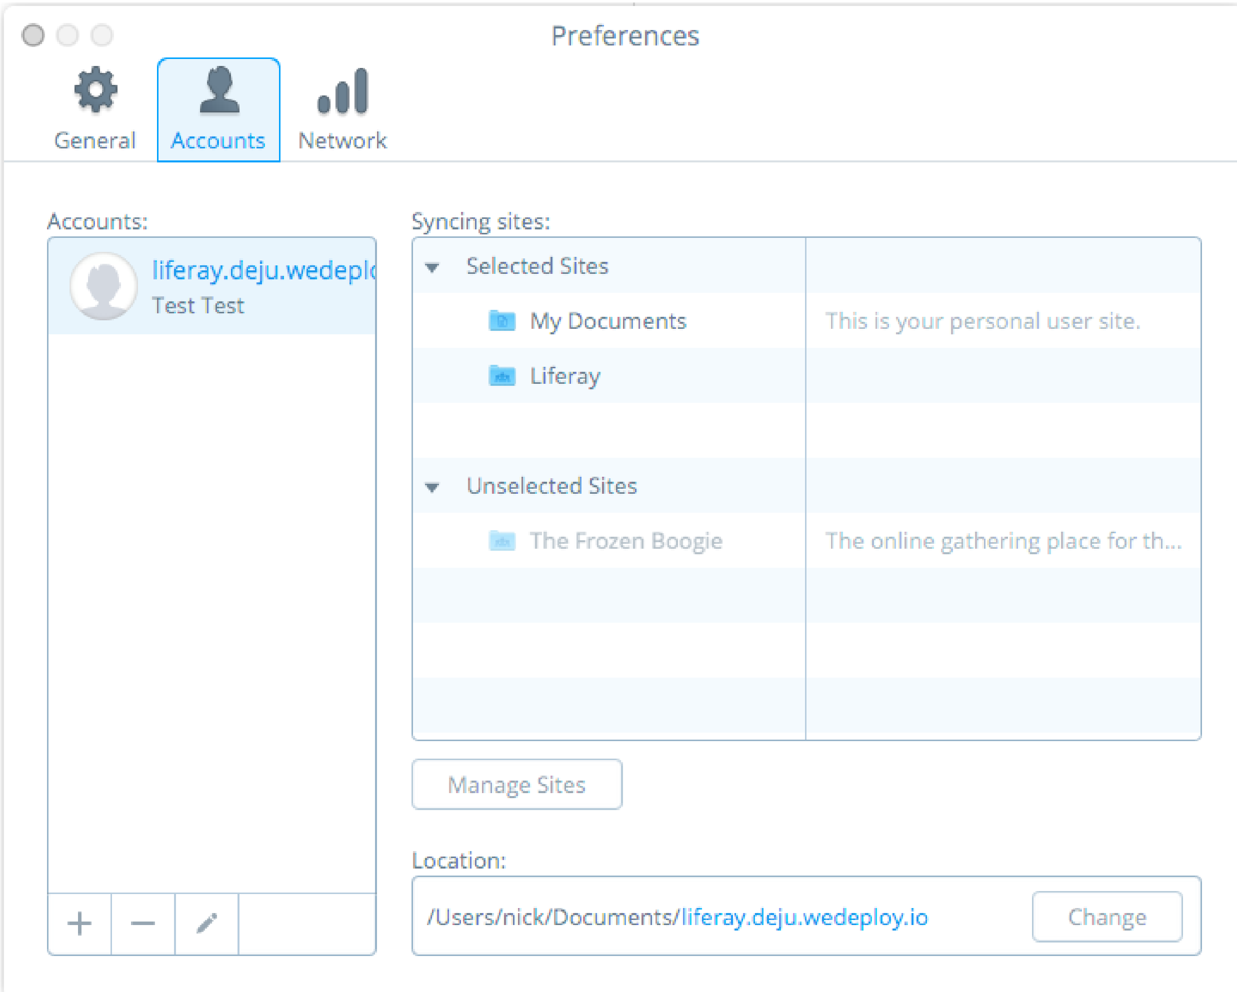

Figure 47.9: The Preferences menu's *Accounts* tab lets you manage syncing with Sites per account.

Launch Liferay Sync on startup: Starts Sync automatically each time your machine starts.

Show desktop notifications: Shows a small notification in the corner of your screen when a synced file changes.

Automatically check for updates: Automatically check for new client versions. You can click the *Check Now* button to check for updates manually.

Finally, the Preferences menu's *Network* tab controls how Sync transfers data with your Liferay DXP servers. It contains the following options:

Download Rate: To limit the rate at which Sync downloads data, select *Limit to* and then specify the rate.

Upload Rate: To limit the rate at which Sync uploads data, select *Limit to* and then specify the rate.

Enable LAN Syncing: Whether to download updates from other desktop clients on the same local network before downloading from the server. This can help reduce server load and increase data transfer speeds. Note that LAN syncing only works when enabled in the Liferay DXP instance by the administrator, and in other clients.

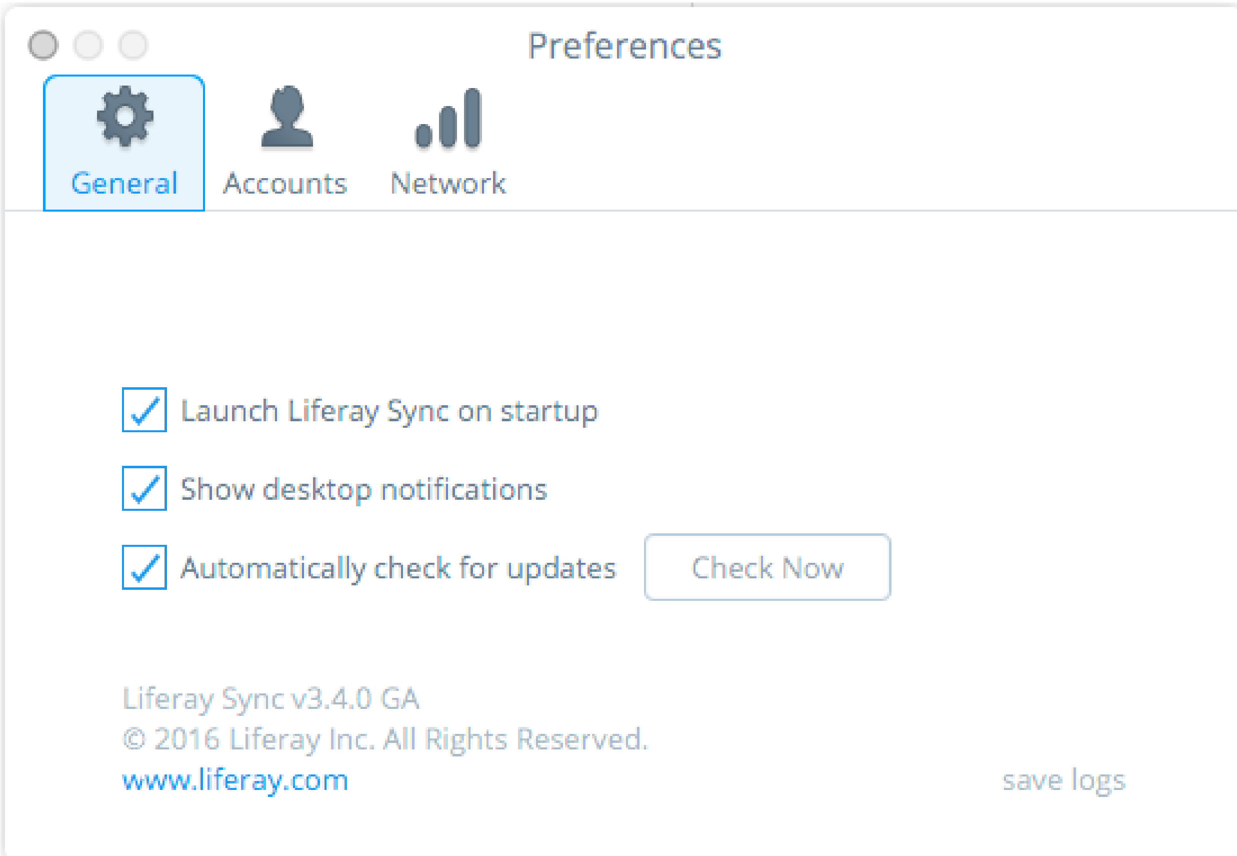

Figure 47.10: The Preferences menu's *General* tab contains settings for Sync's general behavior.

Note that your Liferay DXP administrator can also limit the download/upload rate. In this case, Liferay DXP's settings take precedent. For example, if you set a 5.0 MB/s download rate in the client but Liferay DXP's download limit is 2.0 MB/s, the latter takes precedence. Also, the client's rate applies across all its accounts. For example, if the client connects to three accounts and its download rate is 5.0 MB/s, then the sum of the download rate for all three accounts never exceeds 5.0 MB/s.

47.3 Using Your Local Liferay Sync Folder

Once you configure and run Sync, Sync automatically uploads to Liferay DXP any files you add or modify in your Sync folder. Sync also downloads to your Sync folder any file changes by other users. If you delete a file in your Sync folder, Sync also deletes it from the server and other clients. You should therefore use **extreme caution** when deleting files in your Sync folder. If, however, you accidentally delete a file, all is not lost! The file can still be recovered from the instance's Recycle Bin, which is enabled by default. Note, if the administrator has disabled the Recycle Bin, recovering deleted files is impossible.

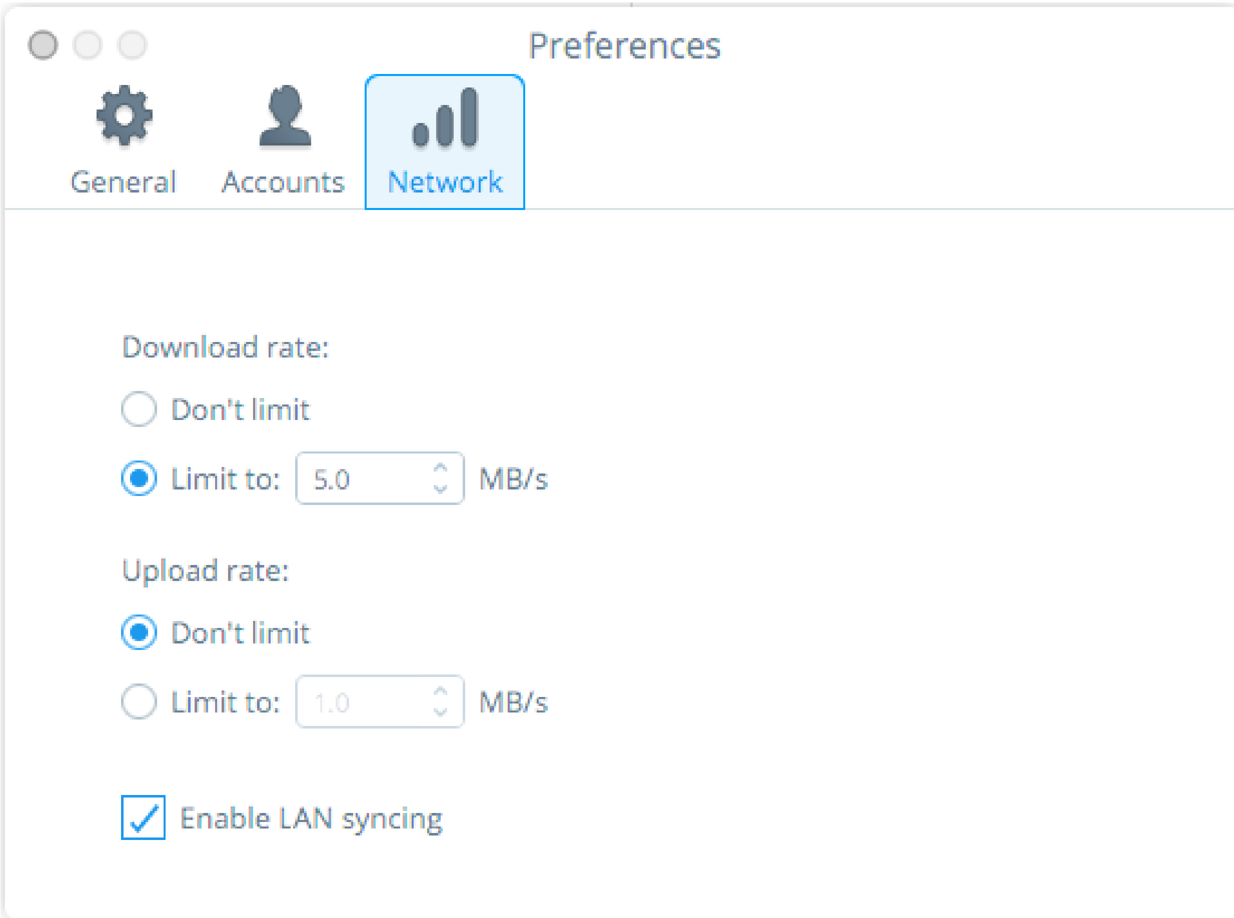

Figure 47.11: The Preferences menu's *Network* tab contains settings for Sync's data transfer behavior.

Warning: Deleting a file in your Sync folder also deletes it in the Liferay DXP instance and in other clients. If you accidentally delete a file, it can be recovered from the instance's Recycle Bin. The Recycle Bin is enabled by default. File recovery is, however, impossible if the instance or Site administrator has disabled the Recycle Bin.

You can run through the following exercise to familiarize yourself with how to create, edit, download, and upload files with Sync. First, open the Sync folder in your file manager and create a new file called `README.txt`. Enter the word `test` in this file. Next, make sure you can access this file in your Site. Go to the Site you want to sync with and navigate to its Documents and Media app. It lists your `README.txt` file.

Download the `README.txt` file to a convenient location on your machine. Open the file and check that it still says `test`. Now open the `README.txt` file in your Sync folder and edit it so that it says `second test`. Once the changes are synced, go back to your browser and refresh the page with your Documents and Media app. Click on the `README.txt` file's name, look at the file information displayed, and check that the file's version number has been incremented.

If you download and open `README.txt` again, it now says `second test`. Your edit was uploaded to the Site! You can be confident that this edit was also downloaded by all other Sync clients connected

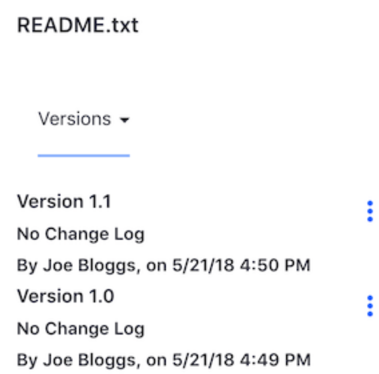

Figure 47.12: Updating a file through Liferay Sync increments the file's version number. You can view a file's version number through the web interface.

to your Site.

Now delete the `README.txt` file from your local Sync folder. When the changes finish syncing, go back to your browser and refresh the page containing your Documents and Media app. The file is gone! The file is also deleted from the local Sync folders of all other Sync clients connected to the Site. Remember this very important rule: deleting files in your local Sync folder deletes them *everywhere*! Next, you'll learn how to use the Sync client for your mobile device.

Chapter 48

Using Liferay Sync on Your Mobile Device

Liferay Sync for Android and iOS contains most of the desktop Sync client's functionality. The mobile client can, however, only be connected to one Liferay DXP account at a time. Also, mobile Sync doesn't automatically download files to your device. To save storage space on your device, the Sync mobile app lets you choose the files you want to work with. As with the Sync desktop clients, the latest versions of Sync on Android and iOS provide a consistent user experience across platforms. While this article details using Sync on Android, the instructions also apply to Sync on iOS.

You need to download and install Sync on your Android or iOS device through its respective app store, the same as you do any other mobile app. To find the app, search Google Play or the App Store for *Liferay*. You can also download Sync from the Liferay Sync downloads page. Once you've installed the Sync app on your device, the rest of the articles in this guide show you how to use it.

48.1 Connecting Liferay Sync Mobile

When Liferay Sync first starts on your mobile device, press the *Get Started* button to begin setup. The setup screen asks for your login credentials and your server's address. Once you enter them, press *Sign In*. After signing in, you see a panel that shows your name, a gear icon for accessing the app's settings, and navigation options *My Sites* and *My Documents*. My Sites and My Documents encompass the Sites in Liferay DXP that you can sync with. My Documents is your personal user Site, while My Sites shows the other Sites with which you can sync. No matter how deep you are in the folder hierarchy of a Site, swiping to the right returns you to this panel. If you're in the first level of My Sites or My Documents, pressing the location bar at the top slides the screen slightly to the right to reveal a compact view of the panel. The following screenshots show both views of the panel.

Press the gear icon to access Sync's settings. Settings shows your account information and an option to sign out of your Liferay DXP instance. Settings also lets you toggle *Security Mode*. Security Mode protects files stored on your device by encrypting them. Using Security Mode requires you to set up a passcode to use when accessing the Sync app. Security Mode protects the files on your device and Liferay DXP instance in the event your device is lost or stolen. You should note, however, downloading and opening files in Security Mode takes slightly longer than usual because

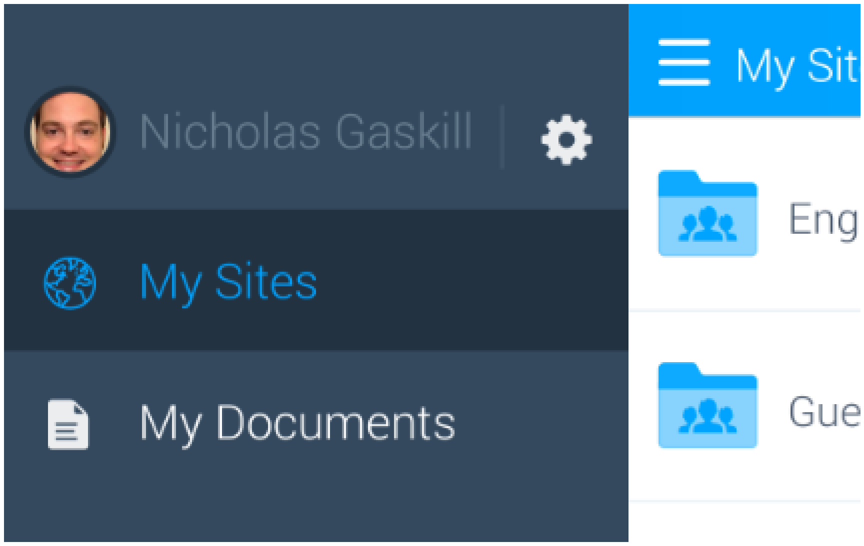

Figure 48.1: This panel lets you access the app's settings, as well as your Sites and documents.

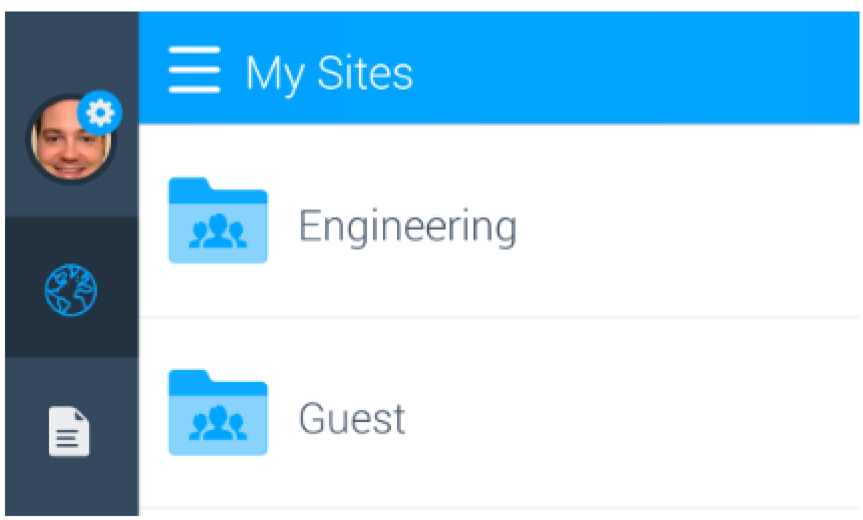

Figure 48.2: Tapping the title bar at the top of My Sites or My Documents opens the main Sync panel's compact view.

the Liferay DXP server must use SSL—if it didn't, your files would be transmitting in the open. Below the Security Mode toggle are the app's version and a link to send app feedback to Liferay.

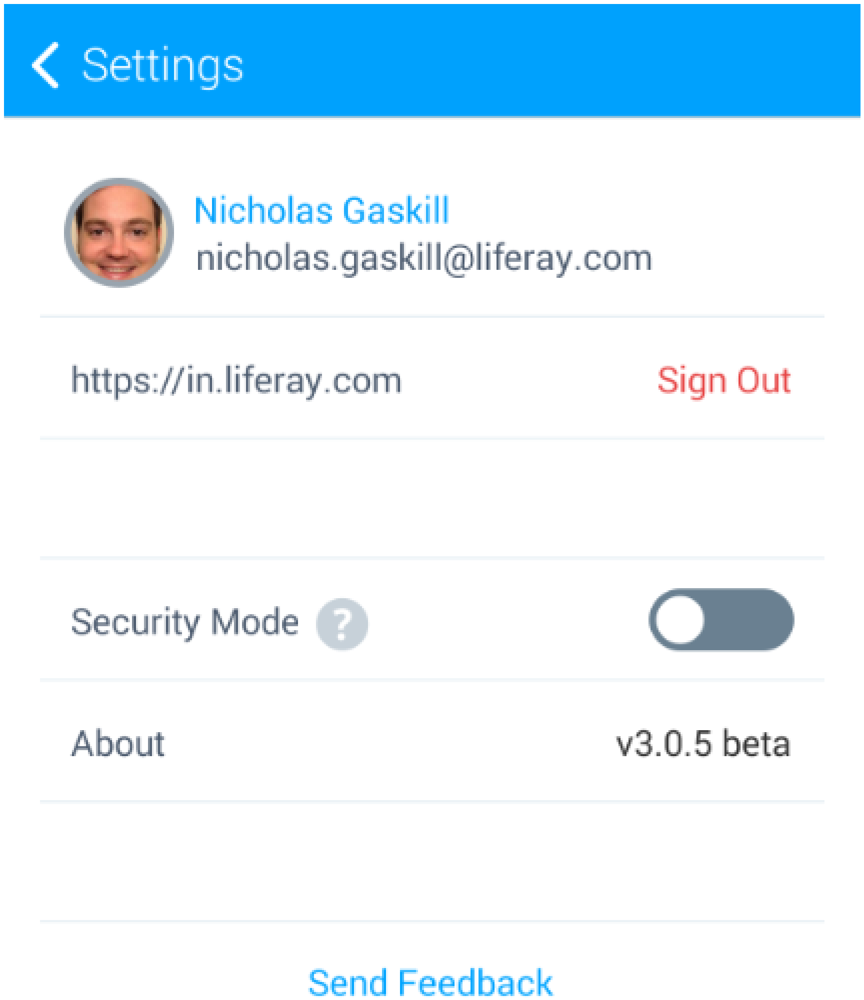

Figure 48.3: The Settings screen for the Sync app lets you sign out of your Liferay DXP instance, enable Security Mode, view the app's version, and send feedback.

48.2 Managing Files and Folders in Liferay Sync Mobile

Whether you're working in My Documents or My Sites, you manage files and folders the same way. Pressing a Site or folder shows you a list of its files and folders. It displays each file's size and modification date. You can refresh the list by pulling down from the top of the screen. Your current location in the navigation hierarchy also appears at the top of the screen alongside a plus icon. Pressing the plus icon launches an upload screen for adding content in the current location. You can add a new folder, upload a file, or launch your device's camera app to take and upload a picture or video. Pressing the *X* icon on the upload screen's top right corner cancels any action and returns you to the current file list.

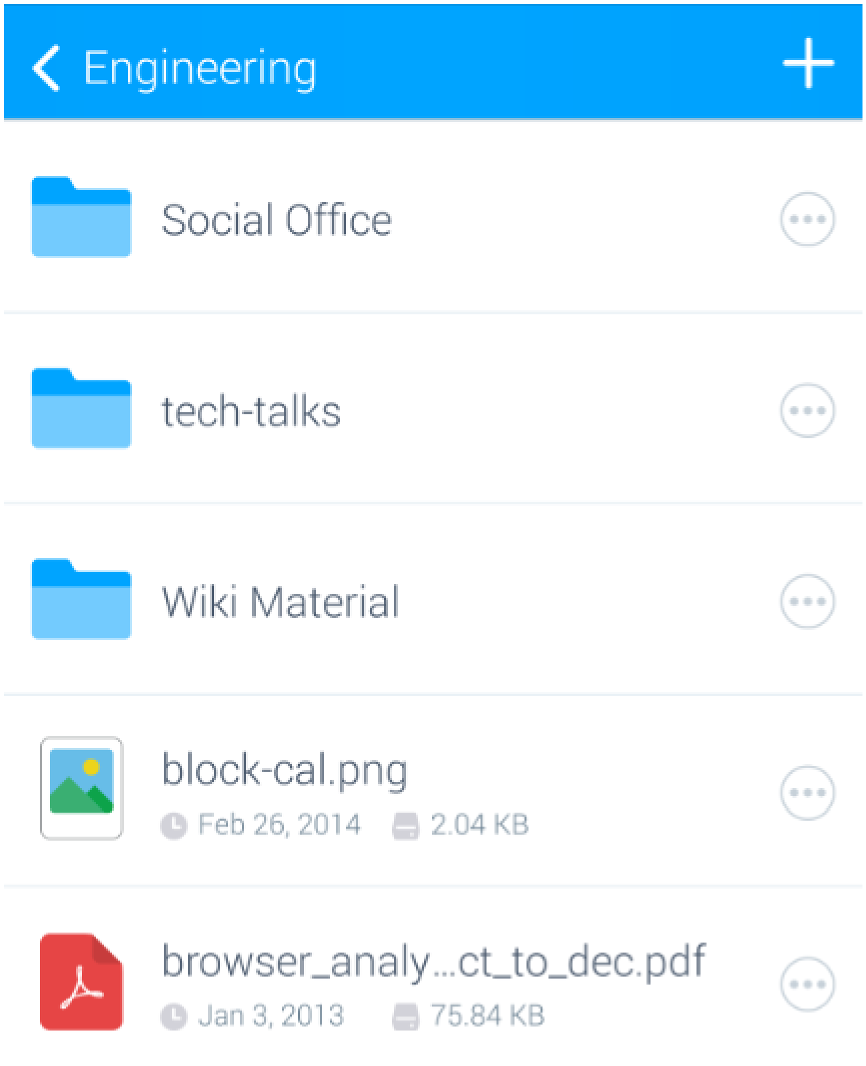

Figure 48.4: Sync shows files and folders in a list.

To download a file to your device, press the file's name in the list. The label that previously showed the file's size and modification date is replaced by a download progress indicator. When the file finishes downloading, your device automatically opens it in the app you've configured to open files of that type. If you haven't configured your device to use a specific app for that file type, you're presented with a list of apps on your device that can open the file. If your device doesn't have an app that can open the file, Sync tells you to install one that can. Downloaded files appear in the list with the file size in blue instead of gray. For example, the screenshot below shows that `LiferayinAction.pdf` is on the device.

The Sync mobile app also lets you move, rename, and delete files and folders. To the right of each file and folder in the list is a circle icon with three dots. Pressing this icon slides open a context menu on the right that lets you move, rename, or delete that item. The screenshots below show these options. Note that you should use **extreme caution** when deleting files or folders. Deleting files or folders in the mobile Sync app also deletes them from Liferay DXP and across any synced clients. Accidentally deleted files can be restored from the Recycle Bin, which is enabled by default.

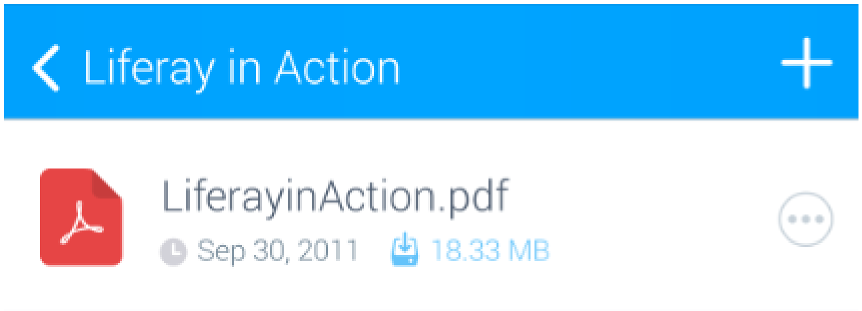

Figure 48.5: Downloaded files appear in the list with their size in blue.

If the instance or Site administrator disables the Recycle Bin, however, recovering deleted files is impossible.

What if you want to delete a file on your device without also deleting it in the instance? Currently, you can only do this by signing out of your account in the app's Settings menu. Doing so removes all downloaded files from your device, but preserves them in the instance. If you're on Android, it may be possible to use a system file browser app to remove downloaded files manually.

Warning: Deleting a file in the mobile Sync app deletes it in Liferay DXP and across any synced clients. If you accidentally delete a file, the instance or Site administrator can restore it from the instance's Recycle Bin. The Recycle Bin is enabled by default. If the instance or Site administrator disables the Recycle Bin, however, recovering deleted files is impossible.

The context menu also provides additional options for files. A small badge on the file icon's top-right corner indicates the file's version in the Liferay DXP instance. You can also use the context menu to share files you've downloaded. Pressing the *Share* icon opens a list of your device's apps capable of sharing the file. To close the context menu and return to the list of files and folders, swipe to the right. The following screenshot shows the options available in a file's context menu.

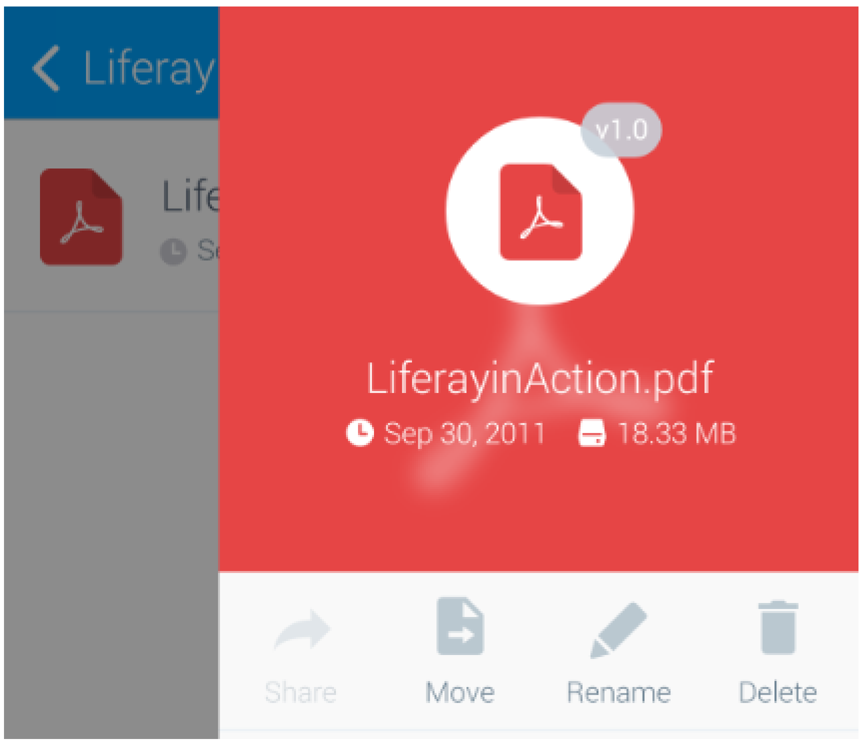

Figure 48.6: The badge on the file's icon shows the file's version in the Liferay DXP instance. You can also share files that you've downloaded.

Chapter 49

Adapting Your Media Across Multiple Devices

Media providers must consider differences between devices (phones, laptops, tablets, etc.) when delivering content: not only their screen size but also their bandwidth and processing capabilities. Liferay DXP's Adaptive Media app allows administrators to control image quality and dynamically adjusts uploaded media to best fit the screen being used.

Note: At this time, Adaptive Media only works for images in blog entries and web content articles.

Adaptive Media integrates with Documents and Media, Blogs, and Web Content. It generates a set of images for use on various screens. When the content is accessed, Adaptive Media checks the screen type and resolution and selects appropriate the appropriate image. Adaptive Media comes preinstalled in Liferay Portal CE 7.1 and Liferay DXP 7.1.

In this section, you'll learn how to manage and use Adaptive Media in Liferay DXP.

49.1 Installing Adaptive Media

The Adaptive Media app is installed in Liferay DXP by default. The following sections describe the Adaptive Media app's modules, and how to prepare Adaptive Media to handle animated GIFs.

Note: Since the Adaptive Media app is installed by default, it's updated via Liferay DXP Fix Packs and Liferay Portal CE GA releases. Using Liferay Marketplace to update the app on Liferay DXP 7.1 causes an error.

Adaptive Media's Modules

Some modules in the Adaptive Media app are mandatory and must be enabled for Adaptive Media to function, while others can be disabled. The Adaptive Media API modules, which export packages for the other modules to consume, are mandatory; disabling one also disables any other modules that depend on it. Here's a list of the Adaptive Media API modules:

- Liferay Adaptive Media API

- Liferay Adaptive Media Content Transformer API
- Liferay Adaptive Media Image API
- Liferay Adaptive Media Image Item Selector API

The Adaptive Media core modules are also mandatory, and must be enabled to ensure that Adaptive Media works as expected:

- Liferay Adaptive Media Document Library
- Liferay Adaptive Media Document Library Item Selector Web
- Liferay Adaptive Media Document Library Web
- Liferay Adaptive Media Image Content Transformer
- Liferay Adaptive Media Image Implementation
- Liferay Adaptive Media Image Item Selector Implementation
- Liferay Adaptive Media Image JS Web
- Liferay Adaptive Media Image Service
- Liferay Adaptive Media Image Taglib
- Liferay Adaptive Media Image Web
- Liferay Adaptive Media Item Selector Upload Web
- Liferay Adaptive Media Web

The Adaptive Media Blogs modules, which ensure that images uploaded to blog entries can be processed and adapted, are optional. Here's a list of these modules:

- Liferay Adaptive Media Blogs Editor Configuration
- Liferay Adaptive Media Blogs Item Selector Web
- Liferay Adaptive Media Blogs Web
- Liferay Adaptive Media Blogs Web Fragment

The Adaptive Media Journal modules are optional. These modules apply Adaptive Media to web content articles:

- Liferay Adaptive Media Journal Editor Configuration
- Liferay Adaptive Media Journal Web

There are two more optional modules included in Adaptive Media:

- **Liferay Adaptive Media Image Content Transformer Backwards Compatibility:** Ensures that content created before the Adaptive Media installation can use adapted images without the need to edit that content manually. It transforms the images both at startup and when a user views the content, which can negatively affect performance. We therefore recommend that you run some performance tests before using this module in production. You can disable this module if you don't have old content, are experiencing performance problems, or your old content doesn't need adapted images.

- **Liferay Adaptive Media Document Library Thumbnails:** Lets thumbnails in Documents and Media use adapted images. For this to work, you must first migrate the original thumbnails to adapted images. We highly recommend that you enable this module, but it's not mandatory.

Great! Now you know the mandatory and optional modules that come with Adaptive Media. The next section discusses the installation requirements for using animated GIFs with Adaptive Media. If you don't need to use GIFs, you can skip ahead to the article on adding image resolutions to Adaptive Media.

Processing Animated GIFs

To process animated GIFs, Adaptive Media uses an external tool called Gifsicle. This tool ensures that the animation works when the GIF is scaled to different resolutions. You must manually install Gifsicle on the server and ensure that it's on the `PATH`. Once it's installed, you must enable it in Adaptive Media's advanced configuration options.

If Gifsicle isn't installed and `image/gif` is included as a supported MIME type in the advanced configuration options, Adaptive Media scales only a GIF's single frame. This results in a static image in place of the animated GIF.

49.2 Adding Image Resolutions

To use Adaptive Media, you must first define the resolutions for the images delivered to users' devices. Adaptive Media then generates new images scaled to fit those resolutions, while maintaining the original aspect ratio.

To access Adaptive Media settings, open the Control Panel and go to *Configuration → Adaptive Media*. Here you can create and manage resolutions.

Note: Adaptive Media configurations apply only to the current Liferay DXP instance.

Once you create a resolution, Adaptive Media automatically generates copies of newly uploaded images in that resolution. Images uploaded before you create the resolution aren't affected and must be adapted separately (see the Generating Missing Image Resolutions section for details).

Adding a New Image Resolution

The number of image resolutions required and the values for each depend on the use case. More resolutions may optimize image delivery, but generating more images requires additional computational resources and storage space. To start, we recommend that you create resolutions to cover common device sizes like mobile phones, tablets, laptops, and desktops. If most users use one device (e.g., all Intranet users have the same company mobile phone), you can create a resolution to target that device.

To add a new resolution, click the *Add* icon (+) on the Adaptive Media configuration page and provide the following information:

- **Name**: The resolution's name (this must be unique). This can be updated if a custom `Identifier` is defined.

- **Max Width**: The generated image's maximum width. If a *Max Height* is given, this field is optional. This value must be at least 1.

- **Maximum Height**: The generated image's maximum height. If a *Max Width* is given, this field is optional. This value must be at least 1.

  Adaptive Media generates images that fit the Max Width and Max Height, while retaining the original aspect ratio. If you only provide one value (either Max Width or Max Height), the generated image scales proportionally to fit within the specified dimension, while maintaining its original aspect ratio. This ensures that adapted images are not distorted.

- **Add a resolution for high density displays (2x):** Defines a scaled up resolution for HIDPI displays. Selecting this option creates a new resolution double the size of the original with the same name and the suffix `-2x`. For example, if the original resolution is `400px` by `300px` (max width by max height), the high density resolution is `800px` by `600px`.
- **Identifier:** The resolution's ID. By default, this is automatically generated from the name. You can specify a custom identifier by selecting the *Custom* option and entering a new *ID*. Third party applications can use this ID to obtain images for the resolution via Adaptive Media's APIs.

Note: Image resolutions and their identifiers can't be updated if the resolution has been used to adapt images. This prevents inconsistencies in generated images.

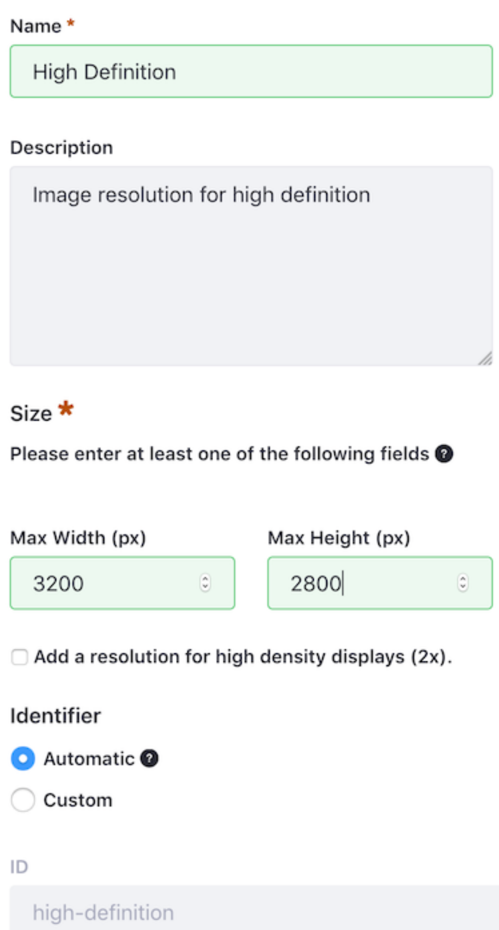

Figure 49.1: The form for adding a new Adaptive Media resolution.

49.3 Managing Image Resolutions

Adaptive Media lets you manage image resolutions and their resulting adapted images. For example, you can disable, enable, edit, and delete resolutions. You can also generate any adapted images that may be missing for a resolution. This article discusses these topics and more.

Disabling Image Resolutions

Disabling an image resolution prevents it from generating adapted images. Any images uploaded after the resolution is disabled use the most appropriate resolution that's still active. Adapted images previously generated by the disabled resolution are still available.

To disable an image resolution, click its *Actions* menu (⋮) and select *Disable*.

Enabling Image Resolutions

Image resolutions are enabled by default. If you need to enable a disabled resolution, click that resolution's *Actions* menu (⋮) and select *Enable*.

While a resolution is disabled, it doesn't generate adapted images for new image uploads. After enabling a resolution, you should generate the adapted images that weren't generated while it was disabled (see the Generating Missing Image Resolutions section for instructions on this).

Editing Image Resolutions

You can't edit an image resolution that already has adapted images. This prevents odd behavior (of the adapted images–you're still free to be as odd as you want). This is because any changes would only be applied to images uploaded after the edit, creating an inconsistent set of adapted images. Odd indeed.

Therefore, editing an image resolution is only possible if Adaptive Media hasn't yet generated adapted images for it. If you must change the values of a resolution that already has adapted images, you must delete that resolution and create a new one with the new values. The next section discusses deleting resolutions.

Deleting Image Resolutions

Be careful when deleting an image resolution, as any adapted images it created are irretrievably lost and are not automatically replaced by new image resolutions you create.

Follow these steps to delete an image resolution:

1. Disable the resolution. You can't delete enabled resolutions. This prevents the accidental deletion of image resolutions.

2. To delete the resolution and all its adapted images, select *Delete* from the resolution's Actions menu (⋮).

Generating Missing Adapted Images

If Adaptive Media hasn't generated all the images you need–say, if new images were uploaded before a new image resolution was created or while the resolution was disabled–you must generate the missing images manually.

To manually generate missing adapted images,

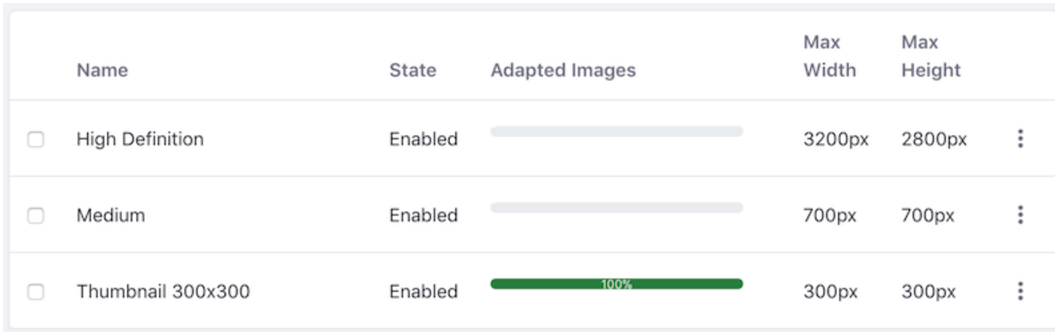

Figure 49.2: The *Adapted Images* column shows the percentage of images that are adapted for each resolution.

1. For a single resolution, select *Adapt Remaining* from the resolution's Actions menu ( ).

2. For all resolutions at once, select *Adapt All Images* from the Actions menu in the Control Menu at the top of the page.

The Recycle Bin and Adapted Images

You can't move adapted images directly to the Recycle Bin. But if the original image is in the Recycle Bin, the corresponding adapted images behave as if they are in the Recycle Bin and users can't view them.

Note: URLs that point to adapted images whose original image is in the Recycle Bin return an error code of `404 Not Found`.

If the original image is restored from the Recycle Bin, the adapted images are accessible again.

Awesome! Now you know how to manage image resolutions in Adaptive Media. Next, you'll learn about creating content with adapted images.

49.4 Creating Content with Adapted Images

Adaptive Media is mostly invisible for blog and web content creators. Once an image is added to the content, the app works behind the scenes to deliver an adapted image appropriate to the device in use. Content creators select an image when adding it to their content—they don't have to (and can't) select an adapted image. Adaptive Media identifies each adapted image in the content's HTML with a `data-fileentryid` attribute that is replaced with the latest adapted image when the user views the content. This lets Adaptive Media deliver the latest adapted images to your content, even if the content existed prior to those images.

Note: If Adaptive Media is uninstalled, the original images are displayed in the blog entries and web content articles.

Including Adapted Images in Content

Since Adaptive Media delivers the adapted images behind the scenes, content creators should add images to blog entries and web content as usual: by clicking the image button in the editor and then selecting the image in the file selector.

However, there are some important caveats. When using the file selector to include an image for a blog entry, Adaptive Media works only with images added from the *Blog Images*, *Documents and Media*, and *Upload* tabs. Additionally, adapted images can only be applied to a blog entry's content–cover images excluded. Adaptive Media works for images added to a blog entry via drag and drop, as the image is automatically uploaded to the Blog Images repository, adapted, and then included in the blog entry's content. You can see this by inspecting the HTML and checking that the image contains the `<img>` tag and `data-fileentryid` attribute.

For web content articles, Adaptive Media works only with images added from the file selector's *Documents and Media* tab. Unlike blogs, Adaptive Media doesn't deliver adapted images for images added to web content articles via drag and drop.

For both blog entries and media content articles, Adaptive Media doesn't work with images added from the file selector's *URL* tab. This is because the image is linked directly from the URL and therefore provides no image file for Adaptive Media to copy.

Note that you can see the `<img>` tag and `data-fileentryid` attribute in the HTML of a blog entry or a web content article while you're writing it. When the content is displayed, the HTML is automatically replaced and looks similar to this:

```
<picture>

    <source media="(max-width:850px)"
    srcset="/o/adaptive-media/image/44147/med/photo.jpeg">

    <source media="(max-width:1200px) and (min-width:850px)"
    srcset="/o/adaptive-media/image/44147/hd/photo.jpeg">

    <source media="(max-width:2000px) and (min-width:1200px)"
    srcset="/o/adaptive-media/image/44147/ultra-hd/photo.jpeg">

    <img src="/documents/20140/0/photo.jpeg/1992-9143-85d2-f72ec1ff77a0">

</picture>
```

This example uses three different images, each with a different resolution. A `source` tag defines each of these images. Also note the original image (`img`) is used as a fallback in case the adapted images aren't available.

Using Adapted Images in Structured Web Content

To use adapted images in structured web content, content creators must manually include an image field in the web content's structure. Then they can reference that image field in the matching template by selecting it on the left side of the editor. Here's an example snippet of an image field named `Imagecrrf` included in a template:

```
<#if Imagecrrf.getData()?? && Imagecrrf.getData() != "">
  <img data-fileentryid="${Imagecrrf.getAttribute("fileEntryId")}"
  alt="${Imagecrrf.getAttribute("alt")}" src="${Imagecrrf.getData()}" />
</#if>
```

This snippet includes the `data-fileentryid` attribute to ensure that Adaptive Media replaces the image with an adapted image. If you inspect the resulting web content's HTML in the editor's code view, you should see a tag like this:

```
<img data-fileentryid="37308"
src="/documents/20143/0/photo.jpeg/85140258-1c9d-89b8-4e45-d79d5e262318?t=1518425" />
```

Note the `<img>` tag with a `data-fileentryid` attribute. Adaptive Media uses the file entry ID to replace the `<img>` element automatically with a `<picture>` element that contains the available adapted images for each resolution (see the `<picture>` example above).

Staging Adapted Images

Adaptive Media is fully integrated with Liferay DXP's content staging and export/import functionality. Adaptive Media includes adapted images in staged content when published, and can update those images to match any new resolutions.

Similarly, when content that contains adapted images is exported, Adaptive Media exports those images in the LAR file. That LAR file can then be imported to restore or transfer that content, along with its adapted images.

Adaptive Media doesn't regenerate adapted images during export/import or the publication of staged content. To improve performance, Adaptive Media instead reuses the existing adapted images.

Awesome! Now you know how create content that contains adapted images. You also know how Adaptive Media includes adapted images in the content's HTML.

49.5 Migrating Documents and Media Thumbnails to Adaptive Media

Liferay DXP automatically generates thumbnails for images in Documents and Media. Once you deploy the Adaptive Media app, however, Liferay DXP doesn't display thumbnails until you migrate them to Adaptive Media. This article walks you through this migration process.

Note: You must be a Portal Administrator to perform the actions described here.

You'll get started by creating image resolutions for the thumbnails in Adaptive Media.

Adding the Replacement Image Resolutions

To migrate the existing Documents and Media thumbnails, you must add new image resolutions in Adaptive Media that have maximum height and maximum width values that match the values specified in the following portal properties:

```
dl.file.entry.thumbnail.max.height

dl.file.entry.thumbnail.max.width

dl.file.entry.thumbnail.custom1.max.height

dl.file.entry.thumbnail.custom1.max.width

dl.file.entry.thumbnail.custom2.max.height

dl.file.entry.thumbnail.custom2.max.width
```

Note: Some of these properties may not be enabled. You need only create image resolutions in Adaptive Media for the enabled properties.

To create the new Image Resolutions, follow the instructions found in the Adding Image Resolutions section of the Adaptive Media user guide.

Now you're ready to to create the Adaptive Media images.

Creating the Adaptive Media Images

Once the required image resolutions exist, you can convert the Documents and Media thumbnails to Adaptive Media images. As mentioned in the Adaptive Media installation guide, the module *Liferay Adaptive Media Document Library Thumbnails* (which is included in the Adaptive Media app) enables this functionality.

There are two different ways to migrate the Documents and Media thumbnails to Adaptive Media:

- **Adapt the images for the thumbnail image resolution:** This scales the existing thumbnails to the values in the Adaptive Media image resolutions, which can take time depending on the number of images. We only recommend this approach when there isn't a large number of thumbnails to process, or if you prefer to generate your images from scratch. This approach is covered in more detail in the Generating Missing Adapted Images section of the Adaptive Media user guide.

- **Execute a migrate process that reuses the existing thumbnails:** This copies the existing thumbnails to Adaptive Media, which performs better because it avoids the computationally expensive scaling operation. The next section describes the steps to run this process.

Running the Migration Process

The migration process is a set of Gogo console commands. You can learn more about using the Gogo console in the Felix Gogo Shell tutorial.

Follow these steps to migrate your thumbnails from the Gogo console:

1. Run the `thumbnails:check` command. For each instance, this lists how many thumbnails are pending migration.

2. Run the `thumbnails:migrate` command. This executes the migration process, which may take a while to finish depending on the number of images.

3. Run the `thumbnails:cleanUp` command. This deletes all the original Documents and Media thumbnails and updates the count returned by `thumbnails:check`. Therefore, you should **only** run `thumbnails:cleanUp` after running the migrate command and ensuring that the migration ran successfully and no images are pending migration.

Note: If you undeploy Adaptive Media at some point after running the migration process, you must regenerate the Documents and Media thumbnails. To do this, navigate to *Control Panel* → *Configuration* → *Server Administration* and click *Execute* next to *Reset preview and thumbnail files for Documents and Media*.

Great! Now you know how to migrate your Documents and Media thumbnails to adapted images.

49.6 Advanced Configuration Options

Adaptive Media's advanced configuration options are available in System Settings. Open the Control Panel and go to *Configuration* → *System Settings*, then select the *Collaboration* tab. There are two configurations for Adaptive Media: *Adaptive Media* and *Adaptive Media Image*.

The *Adaptive Media* configuration affects all virtual instances and is related to Adaptive Media's asynchronous processing. These values can be modified to improve performance for specific scenarios or use cases. The following configuration options are available:

- **Max Processes:** The maximum number of processes for generating adapted media. The default value is `5`.
- **Core Processes:** The number of processes always available for generating adapted media. The default value is `2`. This setting can't exceed the *Max processes* setting.

Warning: Larger values for Max Processes and Core Processes may cause out of memory errors, as processing more images at once can consume large amounts of memory. Out of memory errors can also occur if the source images Adaptive Media uses to generate adapted images are large. You can restrict the maximum size of such images via the *Max Image Size* setting in the *Adaptive Media Image* configuration, which is described next. You should run performance tests to optimize these settings for the amount of memory available on your system.

The *Adaptive Media Image* configuration can be different for each virtual instance. The following configuration options are available:

- **Supported MIME Types:** A list of the image MIME types that Adaptive Media supports. If an image is uploaded and its MIME type isn't in this list, Adaptive Media ignores the image. By default, this list contains many common MIME types.
- **Gifsicle:** To scale animated GIFs, Adaptive Media uses an external tool called Gifsicle. First install Gifsicle on the server, ensure that it's on the `PATH` environment variable, and then click the box next to *Gifsicle Enabled*. If Gifsicle isn't installed and `image/gif` is included as a supported MIME type, Adaptive Media scales only one frame of the GIF, making a static GIF.
- **Max Image Size:** Maximum size of the source images that Adaptive Media can use to generate adapted images. Adaptive Media will not generate adapted images for source images larger than this setting. The default value is 10 MB. To generate adapted images for all source images regardless of size, set this to -1. Since generating adapted images from large source images requires significant amounts of memory, you can specify a lower *Max Image Size* to avoid out of memory errors.

Chapter 50

Collaborating

Your users can leverage a robust suite of collaboration apps to get things done and form extensive communities. These apps provide all the features that you would expect of standalone apps, but these apps are integrated: they share a common look and feel, security model, and architecture because of Liferay's development framework. You can use them in combination with user registration and content management features to build a well-integrated, feature-rich site.

This guide shows you how to administer and use the collaboration apps:

- Blogs
- Message Boards
- Wiki
- Alerts and Announcements
- Knowledge Base
- Bookmarks

Chapter 51

Publishing Blogs

The Blogs app's editor has a complete set of WYSIWYG controls that appear when and where you need them. You can also switch to source mode to edit your content's HTML code. In source mode, you can work with light text on a dark background or dark text on a light background. You can even open a dual-screen HTML editor to see your code rendered in real time.

The Blogs app also contains a powerful set of tools for customizing how your blogs appear. For example, display templates like Abstract or Full Content let you choose how much of a blog post appears on a page. You can leverage the built-in display templates or create your own. You can also add a cover image to each of your blog entries. Let's face it—you might not be able to judge a book by its cover, but a blog post with a nice cover image is more likely to draw readers.

Read on to learn about Liferay's blogging platform.

51.1 Adding Blog Entries

Each site contains a built-in blog instance, so you can add blog entries to it right away. The easiest way to do this is in the Site Administration menu. Follow these steps to add a blog entry in Site Administration:

1. Click the Menu button () to open the product menu. Then expand the menu for your site and select *Content → Blogs*. This takes you to the Blogs app for your site. The *Entries* tab is selected by default, which lists the site's blog entries.

2. Click the *Add* button () to create a new blog entry. This presents the blog entry editor. Note that the same form appears when editing a blog entry.

3. The first input field, *Drag & Drop to Upload,* is for optionally adding a cover image for your entry. By default, an Asset Publisher shows this cover image as part of the blog entry's abstract. You can insert any image you like in this field, either via drag and drop or the *Select File* button. The latter lets you choose an existing image in the blog, an image from Documents and Media, or an image that you upload from your machine.

   If you select an image from Documents and Media, you can make changes to it with the Image Editor. Edits you make are applied automatically to a copy of the image, which you can then use as your cover photo.

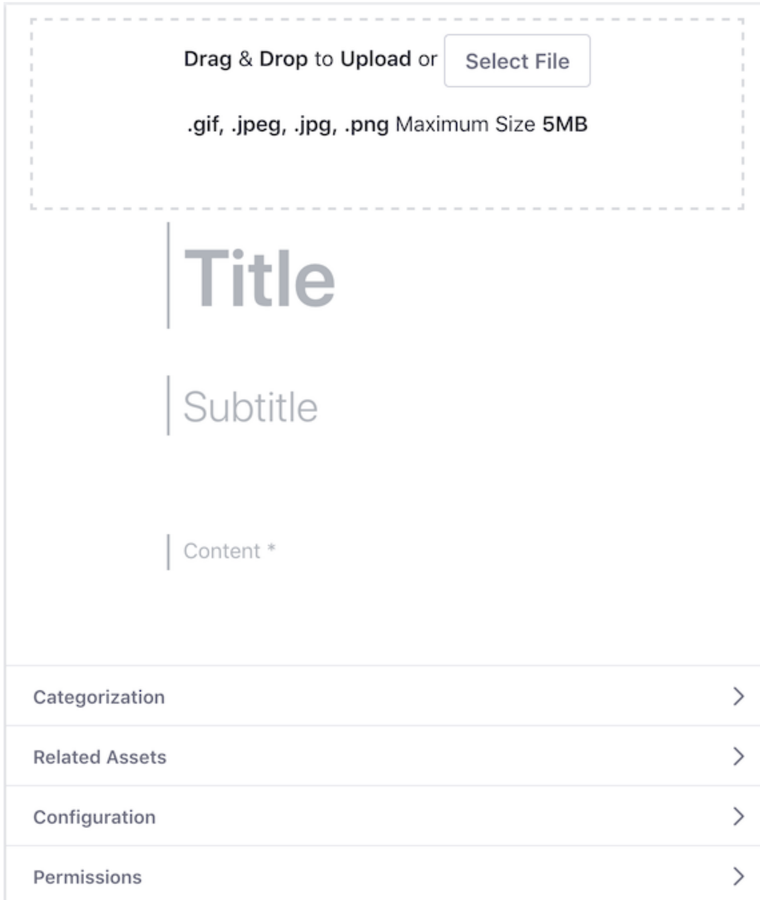

Figure 51.1: This screenshot shows some of the blog entry editor's controls.

Upon upload, the image appears in the pane. To center the image, drag it into place. You can also add a caption. If you want to select a different image, click the *Change* icon ( ). Clicking the trash can icon removes the image from the blog entry.

4. Enter a title for your blog entry. You can also add a subtitle if needed.

5. Enter your blog entry's content in the *Content* field. This field is small at first, but it expands as you add content. The editor displays the editing controls when you need them and hides them from view when you don't. When you select text in your blog post, for example, a bar with context-specific editing controls appears. This keeps your canvas uncluttered so you can focus on writing. You can also add images, videos, and tables in your blog entry's content. See Using the Blog Entry Editor for instructions on creating your blog entry's content.

6. Expand the *Categorization* panel and associate your blog entry with tags and/or categories. Although this is optional, it improves search results for blog entries and gives your users more navigation options. For example, you can add the Tags Navigation app to another column on your blogs page, which lets users browse blog entries by tag.

7. Expand the *Related Assets* panel and choose any other content in your site that you want to associate with this blog entry. Although this is optional, related assets let you tie together

content on your site. For example, you might want to write a blog entry about a discussion that happened on the forums. To link those two assets together, select the forum thread under Related Assets. For more information, see the related assets documentation.

8. Expand the *Configuration* panel if you want to customize your blog entry's URL, abstract, or display date. You can also set whether to allow pingbacks for your blog entry. For the URL, the default selection of *Automatic* generates the URL for you based on the blog entry's title. This URL appears in the *Blog Entry URL* text box. Selecting *Custom* lets you enter your own URL. Note that if you change the blog entry's URL after publishing the entry, the original URL redirects to the new URL.

   You can also specify the blog entry's abstract. Enter a 400 character text-only abstract, or a custom abstract that contains a thumbnail image and a manually written description. The *Small Image* section lets you add a small image that appears when blog entries are displayed in list view. Below the abstract section, you can set the entry's display date and time.

   Note that if you're editing an existing blog entry, the *Send Email Entry Updated* toggle appears. Setting this to *YES* sends an email to any subscribers when the blog entry is updated. You can customize this email when configuring the Blogs app.

   Finally, you can allow *pingbacks* for the blog entry. Pingbacks are XML-RPC requests that are sent automatically when you link to another site. If you link to another site in your blog entry, Liferay DXP sends a pingback to the other site to notify that site that you linked to it. Similarly, if someone links to your blog entry, Liferay DXP can receive a pingback from that site and record the link.

9. Expand the *Permissions* panel to customize your blog entry's permissions. Use the *Viewable by* selector to set who can view the blog entry:

   - **Anyone (Guest Role):** Anyone, including guests, can view the entry.
   - **Site members:** Only site members can view the entry.
   - **Owner:** Only the entry's owner can view the entry.

   Click the *More Options* link to bring up a permissions table that lets you grant or revoke the following permissions for guests and site members:

   - **Update**: Edit and modify the blog entry.
   - **Permissions**: View and modify the blog entry's permissions.
   - **Delete**: Move the blog entry to the Recycle Bin.
   - **View**: View the blog entry.
   - **Update Discussion**: Edit another user's comment on the blog entry.
   - **Delete Discussion**: Delete any comments on the blog entry.
   - **Add Discussion**: Comment on the blog entry.

10. Click *Publish* to publish your blog entry. It now appears in the *Entries* tab.

51.2 Using the Blog Entry Editor

When you create a new blog entry, you create its content with the blog entry editor. This editor is simple yet powerful. Its editing tools are context-aware, only appearing when you need them. They

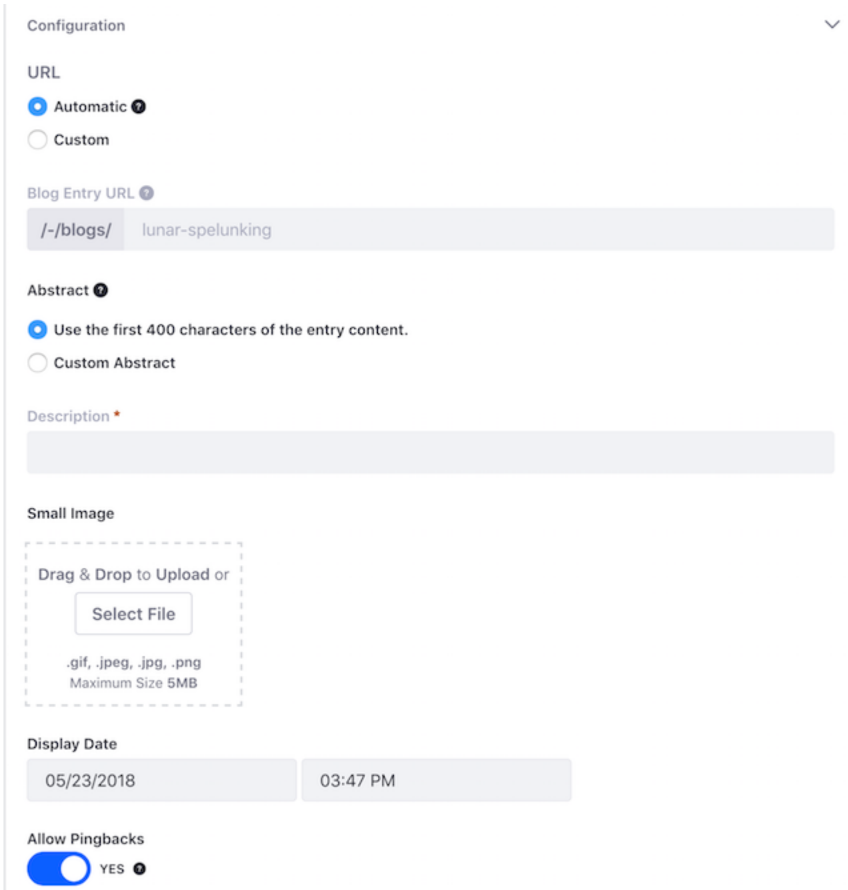

Figure 51.2: When creating a blog entry, the Configuration panel lets you control when and where the blog entry appears, and what to use for the entry's abstract.

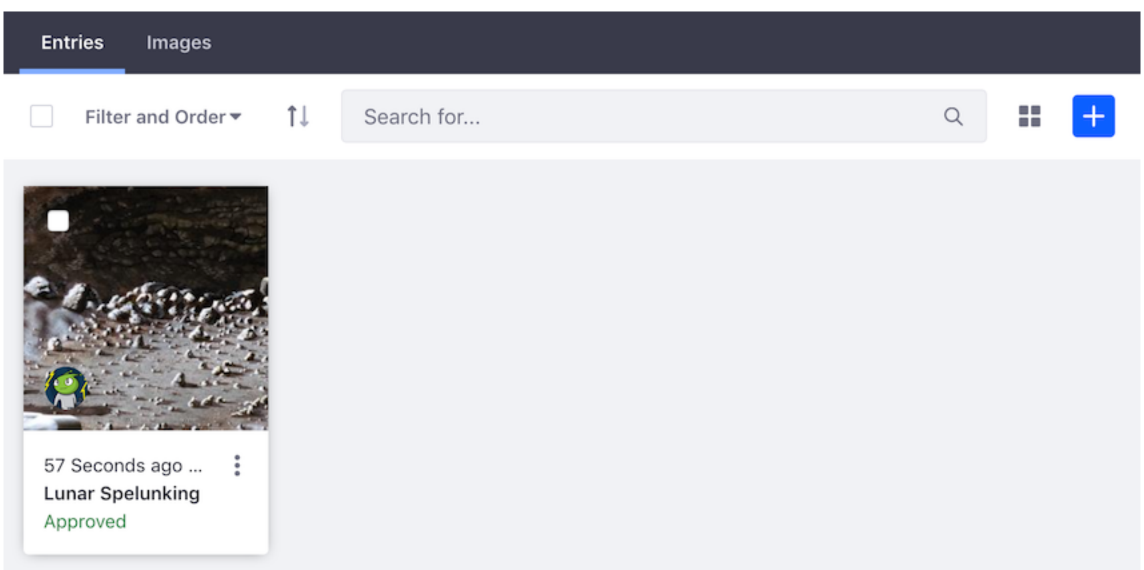

Figure 51.3: The Blogs app in Site Administration lists the site's blog entries.

aren't visible while you're writing, which lets you focus on the task at hand. But when you select text, for example, a toolbar appears that contains buttons for styling the text. There are similar contextual options for adding and editing images, tables, and other types of content. And this is all in the editor's text view. You can switch to code view to edit the content's HTML. Regardless of which view you use, your entry is automatically saved as a draft every 25 seconds, so a browser crash or network interruption won't cause you to lose your work.

This guide shows you how to use this editor to create and edit blog entries.

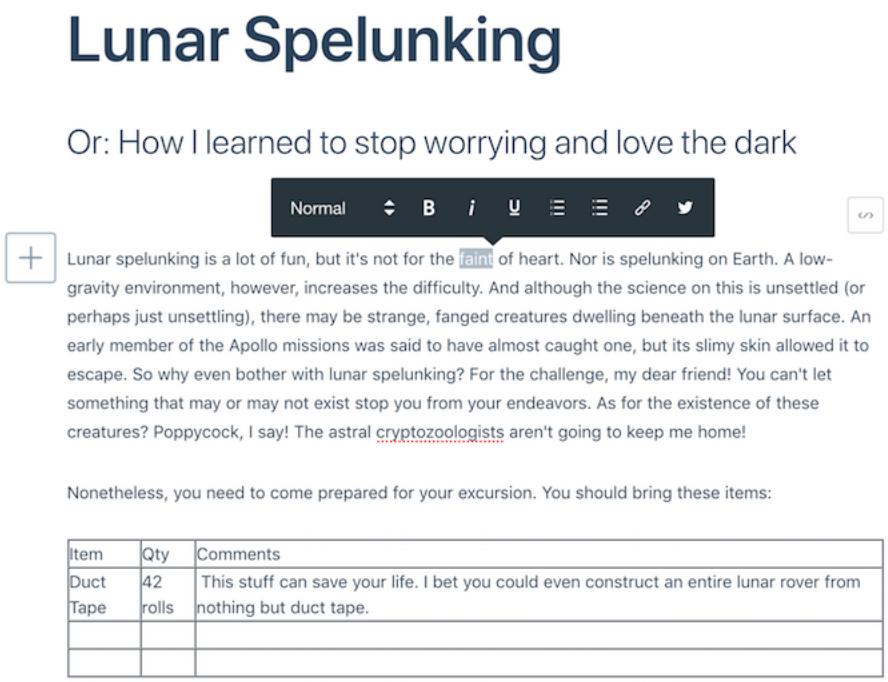

Figure 51.4: This screenshot shows some of the blog entry editor's controls.

Using the Editor's Text View

When you create or edit a blog entry, the editor is in text view by default. Text view is a WYSIWYG editor that lets you enter and edit text and other types of content. To enter text, place your cursor in any text field (e.g., Title, Subtitle, Content, etc.) and type or paste your text. Note that the Content area expands to fit its contents.

To style or format text in the Content area, first select the text. A toolbar appears above the text that contains the following options, as shown in the above screenshot:

- **Text Style:** This selector menu, set to Normal by default, lets you choose the text's style. Normal is typical body text, but you can also select from different heading styles, alert or error message styles, code style, and more.
- **Typeface:** Select bold, italic, or underline.
- **List Style:** Select a numbered or bulleted list.
- **Link:** Link the selected text to a specific URL, or to an item in the portal (e.g. a file in Documents and Media).

- **Twitter:** Generates a link to tweet the selected text.

When you park your cursor in the entry's content area, the *Add* icon (+) appears. If you click this icon, it shows controls for inserting an image, video, table, or horizontal line ( ). Follow these instructions to insert each:

- **Image:** Click the mountain icon, then select or upload an image in the image file selector screen that appears. Alternatively, you can drag-and-drop image files into the content area. You can also use the built-in Image Editor to apply simple edits to an image. Any edits you make are automatically applied to a copy of the image. After you add an image to the blog entry, clicking the image brings up controls for justifying it to the right or left side of the article.

- **Video:** Click the play icon and insert the video's link. The video then appears in your content.

- **Table:** Click the table (grid) icon and then choose the number of rows and columns in your table. When you click inside the table, table editing controls appear. They let you designate the first row and/or column as table headers. The controls also enable you to add rows, columns, and cells. As you type in a cell, the column width automatically adjusts to fit the content.

- **Line**: Click the line icon. A simple, lightweight horizontal line then appears in your content. Such lines are good for separating sections of content in a large blog entry.

Using the Editor's Code View

To switch to code view, click the the *Source* icon (</>) that appears when you place your cursor in the Content area. The following buttons exist at the top-right of code view:

- **Text View** ( ): Switch back to text view.
- **Dark/Light Theme** (/): Switch the code editor between dark and light theme.
- **Fullscreen** ( ): Work in a dual-pane view that shows your HTML code on the left and a preview pane on the right. In this view, you can arrange the HTML and preview panes horizontally or vertically. You can also hide the preview pane so the HTML editor takes up the entire window space.

51.3 Managing Blog Entries

The Blogs app in Site Administration helps bloggers and blog administrators manage blog entries. To access this app, click the Menu button ( ) to open the product menu, then expand the menu for your site and select *Content → Blogs*. The *Entries* tab is selected by default, which lists the site's blog entries.

You can use the Management Bar to manage your site's blog entries. If you've added blog entries via your site's Blogs app, then you're already familiar with the Management Bar's *Add* button ( ). The sections that follow describe the rest of the Management Bar.

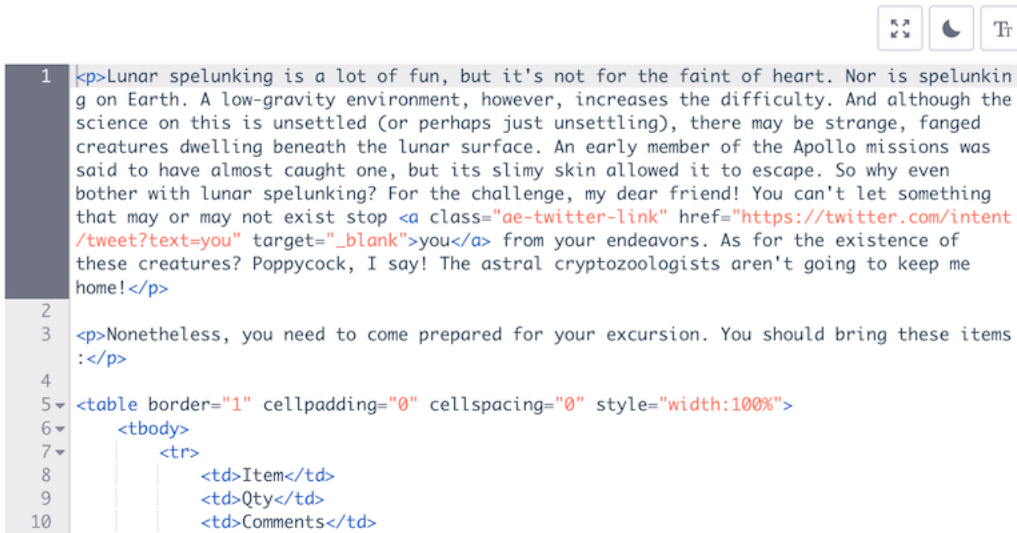

Figure 51.5: Editing in code view lets you work with your blog entry's underlying HTML.

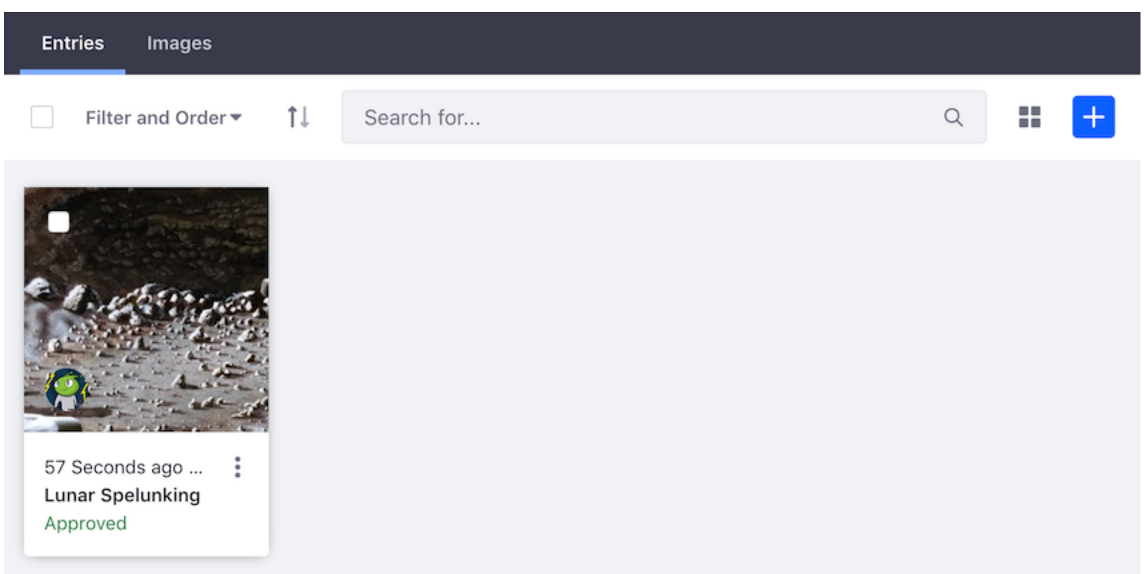

Figure 51.6: The Blogs app in Site Administration lists the site's blog entries.

View Types

The *View Types* button is to the left of the Add button. It lets you choose how to display the blog entries in the Blogs app. The View Types button's icon depends on the selected view type:

- **Cards** (): Shows a card-like rendering of the blog entry, with the author's profile picture. If the entry doesn't contain a cover image, a generic rendering of the entry is displayed. Each card also contains the entry's timestamp, title, workflow status (e.g., Approved, Draft, etc.), and an Actions menu ( ).

- **List** (): Shows the same information as the Cards view type, in a list with the author's profile picture instead of the blog entry's cover image.

- **Table** (): Shows the same information as the other view types, in a list with no file renderings. Also, the blog entry information is in columns.

Finding and Arranging Blog Entries

The Management Bar also contains tools that help you locate and arrange blog entries. The most prominent of these tools is the *Search* bar, where you can find files by keywords.

To the left of the Search bar, the Sort button () arranges entries in ascending or descending order.

You can also arrange entries via the *Filter and Order* selector using these criteria:

- **All:** Shows all of the site's entries.
- **Mine:** Shows only the current user's entries.
- **Title:** Orders the entries by title.
- **Display Date:** Orders the entries by display date.

Selecting Blog Entries

The checkbox on the left-most side of the Management Bar selects all currently displayed blog entries. Selecting multiple entries lets you act on all of them at once. You can also select multiple entries individually by using the checkboxes for each. When you select one or more entries, the Management Bar changes to reflect the actions you can take on the selected entries.

Click the *Trash* button () to move the selected entries to the Recycle Bin. Click *Deselect all* to deselect all the selected entries and return the Management Bar to its normal view.

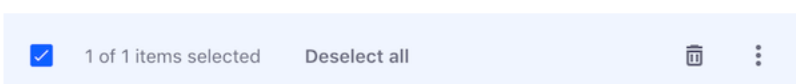

Figure 51.7: With multiple blog entries selected, the management bar changes to reflect the actions you can take on the selected entries.

51.4 Configuring the Blogs App

By configuring the Blogs app in Site Administration, you can control how the app behaves for all blogs in your site. To access this app, click the Menu button () to open the product menu, then expand the menu for your site and select *Content* → *Blogs*. The *Options* menu () at the top-right of the Blogs app lets you configure permissions and notifications, or import/export the app's content.

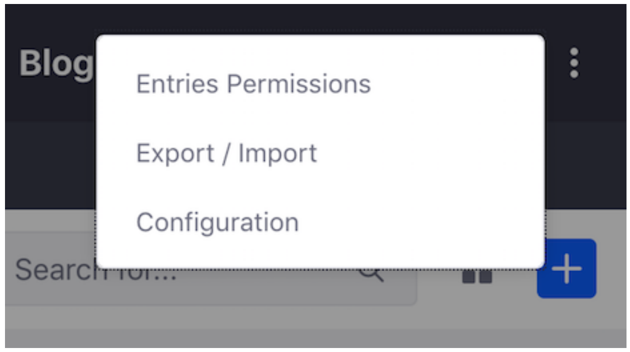

Figure 51.8: You can configure the options for your site's Blogs app.

Here are each of the options available in this menu:

- **Entries Permissions:** Configure the permissions that can be applied to the Blogs app. You can control which roles can add an entry, configure entry permissions, and subscribe to entries.
- **Export/Import:** Export or import a LAR file that contains the Blogs app's content.
- **Configuration:** Configure the following options for the Blogs app, in these tabs:
  - **Email From:** Define the *From* field in the email messages that users receive from Blogs.
  - **Entry Added Email:** Define a subject and body for the emails sent when a new blog entry has been added.
  - **Entry Updated Email:** Define a subject and body for the emails sent when a new blog entry has been updated.
  - **RSS:** Lets you enable RSS subscription and choose how blogs are displayed to RSS readers. The *Maximum Items to Display* selector lets you choose the total number of RSS feed entries to display on the initial page. You can choose up to one hundred to be displayed. The *Display Style* selector lets you choose between *Full Content*, *Abstract*, and *Title* for the entry display in the RSS feed. Lastly, the *Format* selector lets you choose which format the RSS feed uses to deliver the entries: Atom 1.0, RSS 1.0, or RSS 2.0.

51.5 Displaying Blogs

The Blogs app in Site Administration lets you add, manage, and configure your site's blogs. You can then display those blogs by adding a separate Blogs widget to a page. Adding the Blogs widget to a site page creates a shared blog for site members. Adding the widget to a user's personal site (dashboard) creates a blog just for that user. The widget works the same way in both cases. And of course, you can scope a blog to a page to produce a blog instance for just that page.

To add a Blogs widget to a page:

1. Navigate to the page.

2. From the *Add* menu (), open *Widgets* → *Collaboration*.

3. Drag a *Blogs* widget onto the page.

By default, the Blogs widget lists abstracts of the site's recent blog entries. The listing shows each entry's cover image prominently. Each abstract in the listing also shows the number of comments, thumbs up/down ratings, and links to share the entry on Twitter, Facebook, LinkedIn, and other social networking sites. Clicking a blog entry lets you view its full content, where you can also comment on the entry.

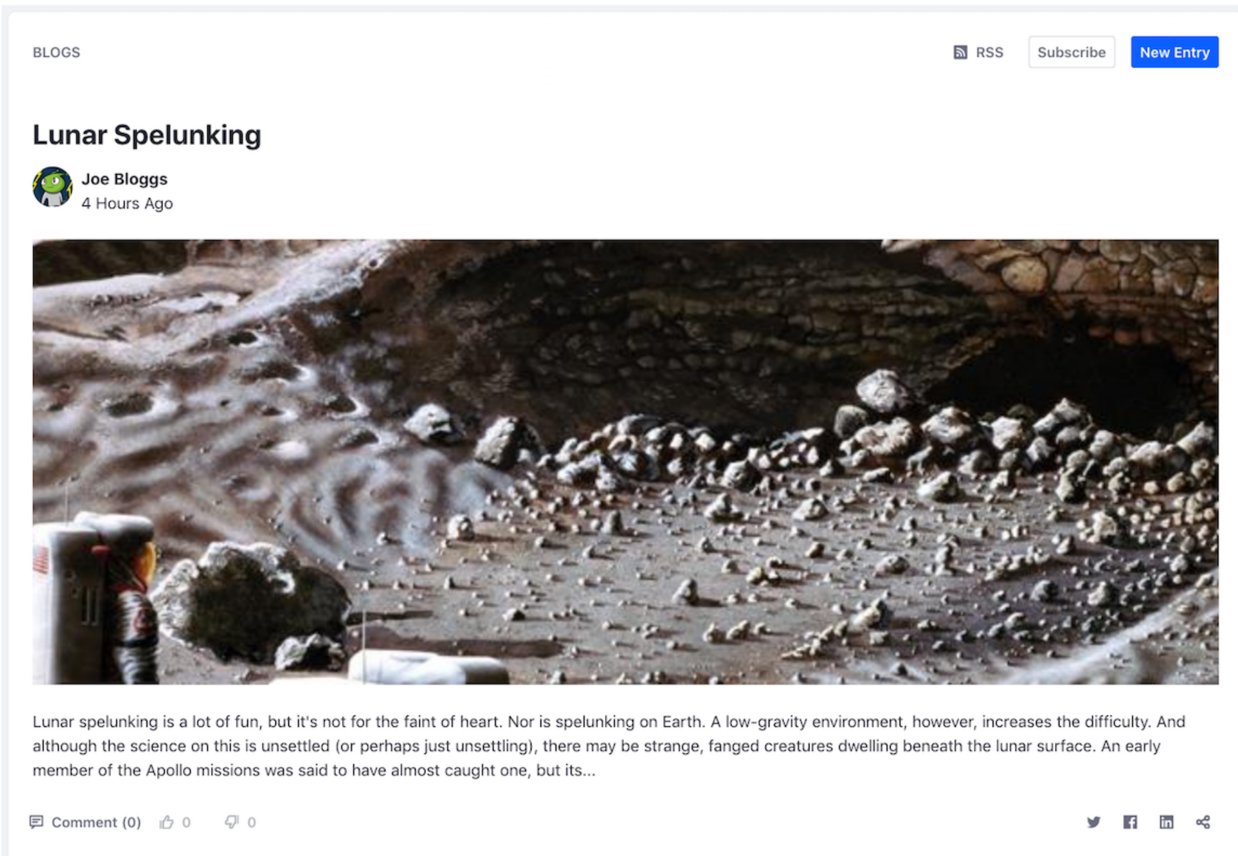

Figure 51.9: Fancy a lunar spelunking trip? This blog entry's abstract lets you know what you're getting into.

There are several display options that let you configure the listing to look the way you want. To configure the widget, click the *Options* icon ( ) in its title bar and select *Configuration*. The display settings are in the *Setup* tab:

- **Enable Ratings:** Whether readers can rate blog entries.

- **Enable Comments:** Whether readers can comment on blog entries.

- **Show View Count:** Whether to show the number views for each entry.

- **Social Bookmarks:** The social networking sites that users can share blog entries with. Only those in the *Current* column are displayed via the share buttons on each blog entry. To move social networking sites between the *Current* and *Available* columns, select the sites and use the arrows between those columns. Similarly, use the up/down arrows beneath the *Current* column to reorder the sites as they appear on each blog entry. For more information, see the social bookmarks documentation.

- **Display Style:** The display style for social bookmarks. *Inline* is the default and displays the social bookmark icons in a row. *Menu* hides them inside a single share menu.

- **Maximum Items to Display:** The total number of blog entries to display on the initial page. You can select up to 75 to display at once.

- **Display Template:** The overall appearance of blog entries in the listing. *Abstract* is the default, and is shown in the above screenshot. You can also choose the following:

  - **Full Content:** Displays each blog entry's full content.
  - **Title:** Displays only the blog entry's title.
  - **Basic:** A stripped-down version of the Abstract, with less text and no cover image.
  - **Card:** Displays each blog entry in a card-like rectangle that shows the cover image, title, author, post date, and a few lines of text.

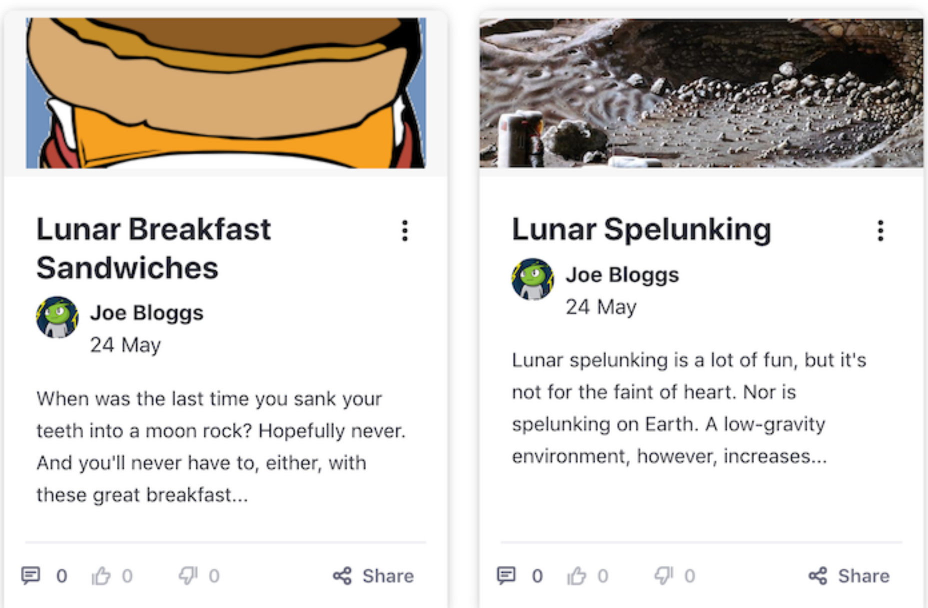

Figure 51.10: The *Card* display template makes your blog posts look like fun little trading cards.

  To select a different application display template (ADT) or create your own, click *Manage Templates*. For more information, see the documentation on application display templates.

- **Enable Report Inappropriate Content:** Whether to let users flag content as inappropriate, which sends an email to administrators.

- **Enable Ratings for Comments:** Whether to let readers rate blog entry comments.

- **Show Related Assets:** Whether to display related content from other apps/widgets in blog entries.

There are also other tabs in *Configuration*:

- **Communication:** Lists public render parameters the widget publishes to other widgets on the page. Other apps/widgets can read and take actions on these. For each shared parameter, you can specify whether to allow communication using the parameter and select which incoming parameter can populate it.
- **Sharing:** Embed the widget instance as a widget on any website, Facebook, Netvibes, or as an OpenSocial Gadget.
- **Scope:** Specify the blog instance the widget displays: the current site's blog (default), the global blog, or the page's blog. If the page doesn't already have a blog instance, you can select scope option *[Page Name] (Create New)* to create a page-scoped blog instance.

When you finish setting the options, click *Save* and then close the dialog box.

51.6 Aggregating Blogs

You can set up a whole web site devoted just to blogging if you wish. The Blogs Aggregator lets you publish entries from multiple bloggers on one page, giving further visibility to blog entries. You can add it to a page from the *Collaboration* category in the *Add* () → *Widgets* menu.

If you click *Configuration* from the *Options* icon ( ) in the widget's title bar, the Blogs Aggregator's configuration page appears. The *Setup* tab contains these options:

- **Selection Method:** Set how the widget selects blogs for display. You can choose *Users* or *Scope*. If you select Users, the widget aggregates the entries of every blogger on your system. To refine the aggregation, you can select an organization by which to filter the users. If you select Scope, the widget aggregates the entries of users in the current scope. This limits the entries to members of the site where the widget resides.
- **Organization:** The organization whose blogs you want to aggregate.
- **Display Style:** Select the overall appearance for blog entries in the widget: *Body and Image*, *Body*, *Abstract* (default), *Abstract without Title*, *Quote*, *Quote without Title*, and *Title*.
- **Maximum Items to Display:** The maximum number of entries the widget displays.
- **Enable RSS Subscription:** Whether to enable an RSS feed of the aggregated entries. This lets users subscribe to an aggregate feed of all your bloggers. Below this option, you can configure how you want to display the RSS feed:
  - **Maximum Items to Display:** The maximum number of RSS items to display.
  - **Display Style:** The overall appearance of each entry in the RSS feed: *Abstract*, *Full Content*, or *Title*.
  - **Format:** The language to use for your RSS feed: *Atom 1.0*, *RSS 1.0*, or *RSS 2.0*.
  - **Show Tags:** Whether to display each entry's tags.

Here are descriptions for the other tabs in the Blogs Aggregator's configuration:

- **Sharing:** Embed the widget instance as a widget on any website, Facebook, Netvibes, or as an OpenSocial Gadget.
- **Scope:** Specify the blog instance the widget displays: the current site's blog (default), the global blog, or the page's blog. If the page doesn't already have a blog instance, you can select scope option *[Page Name] (Create New)* to create a page-scoped blog instance.

When you finish setting the options, click *Save* and then close the dialog box. You'll notice that the Blogs Aggregator looks very much like the Blogs widget, except that it shows entries from multiple blogs.

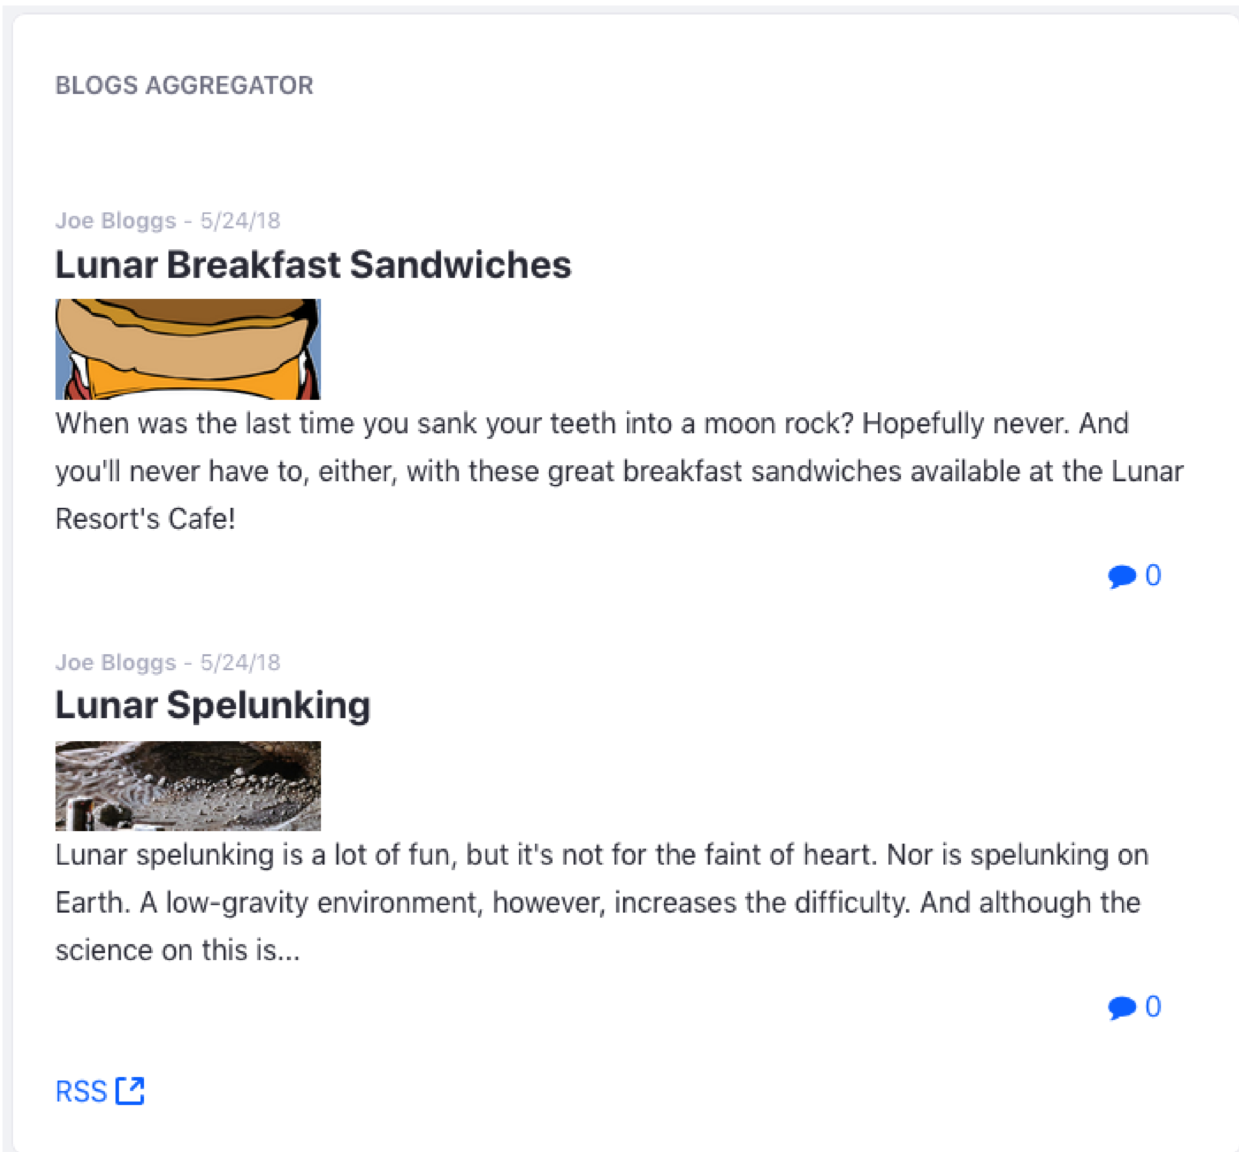

Figure 51.11: The Blogs Aggregator lets you display blog entries authored by multiple authors from different sites.

51.7 Highlighting Recent Bloggers

The Recent Bloggers widget lets you highlight the work of your site's most recent blog authors. This widget lists each recent author's name, profile picture, and number of posts. You can add the Recent Bloggers widget to a page from the *Collaboration* category in the *Add* () → *Widgets* menu.

To access the widget's configuration options, click *Configuration* from the *Options* menu ( ) in the widget's title bar. The *Setup* tab appears first:

- **Selection Method:** Set how the widget selects blogs authors to highlight. You can choose *Users* or *Scope*. If you select Users, the widget aggregates every recent blogger on your system. To refine the aggregation, you can select an organization by which to filter the users. If you select Scope, the widget aggregates the recent bloggers in the current scope. This limits the entries to members of the site where the widget resides.
- **Organization:** The organization whose recent bloggers you want to aggregate.
- **Display Style:** Select how the widget displays recent bloggers: *User Name and Image*, or *User*.
- **Maximum Bloggers to Display:** Select the maximum number of recent bloggers the widget displays.

Here are descriptions for the other tabs in the widget's configuration:

- **Sharing:** Embed the widget instance as a widget on any website, Facebook, Netvibes, or as an OpenSocial Gadget.
- **Scope:** Specify the blog instance the widget displays: the current site's blog (default), the global blog, or the page's blog. If the page doesn't already have a blog instance, you can select scope option *[Page Name] (Create New)* to create a page-scoped blog instance.

When you're finished setting the options, click *Save*. Then close the dialog box.

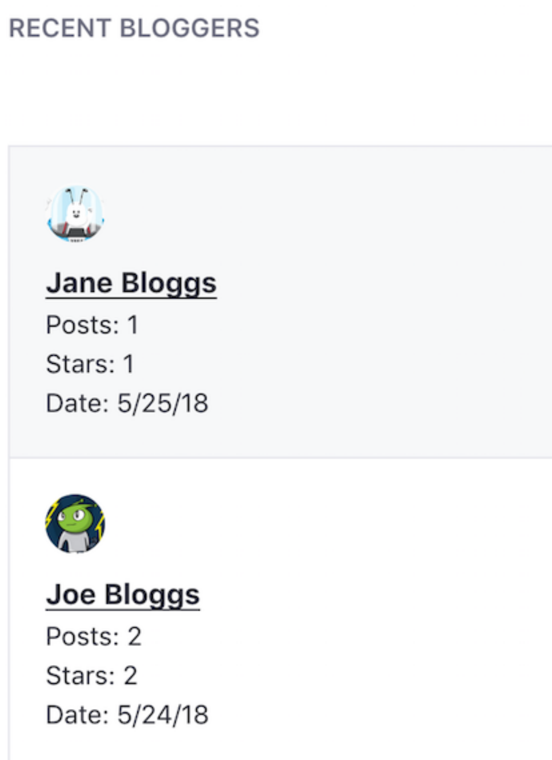

Figure 51.12: You can show off your site or organization's most recent bloggers from the Recent Bloggers app.

CHAPTER 52

CREATING FORUMS WITH MESSAGE BOARDS

Although you're likely already familiar with what a modern forum can do, here's a sampling of what users and administrators can do with Liferay DXP's Message Boards app.

Users can:

- Start and reply to threads
- Mark a thread as a question and select an answer from the replies.
- Subscribe to threads
- Author posts in BBCode or with the standard WYSIWYG editor
- Assign thread priority (e.g., sticky, announcement, etc.)
- Attach files to a thread
- Rate a thread (e.g., like/dislike)
- And more

Administrators can:

- Organize threads into categories and subcategories
- Scope a message board to a page, a Site, or the entire portal.
- Publish threads via RSS
- Rank users by the number of messages they post and assign labels to these rankings (e.g., novice, legend, etc.)
- Create and modify thread priorities (e.g., sticky, announcement, etc.)
- Configure email notifications for thread activity
- And more

As you can see, there's something for everyone!

The Message Boards app also integrates with the rest of Liferay DXP's features. In many web sites, it's obvious that there's no link between the main Site and the message boards. In some cases, users are even required to register twice: once for the main Site and once for the message boards. Sometimes it's even three times: once for the Site, once for the message boards, and once for the shopping cart. By providing a message boards app along with all of the other apps and widgets, Liferay DXP provides a unique, integrated approach to building Sites. Administrators can concentrate on building their Site while the integration work is done for them.

52.1 Creating Message Boards

You can create and manage message boards in the Global, Site, and Page scopes. Regardless of scope, you manage a message board via the Site Administration menu. The following sections show you how to use this menu to manage message boards in each of these scopes.

Site-scoped Message Boards

By default, the Message Boards app in Site Administration is scoped to the current Site. To administer this message board, open the *Menu* (), expand the menu for your Site, then navigate to *Content* → *Message Boards*. The Message Boards administration screen then appears. Note that the options available on this screen are the same regardless of scope. The next sections show you how to change scope and then access this screen.

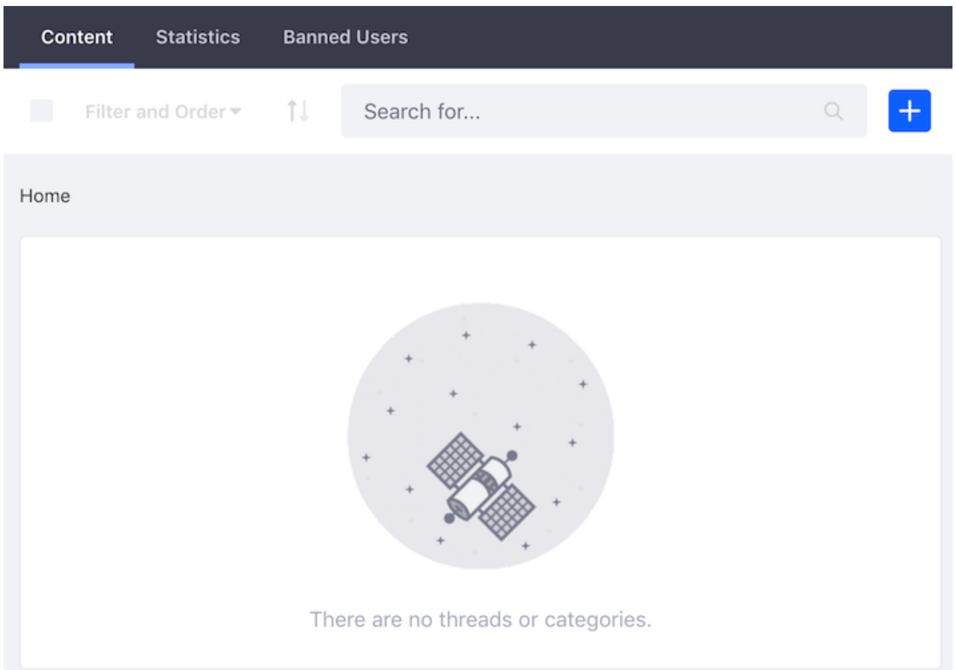

Figure 52.1: A Message Board instance starts empty, ready for you to configure for your purposes.

Page-scoped Message Boards

If you need a page-scoped message board, you must add a Message Boards widget to that page and then set its scope to the page. Follow these steps:

1. Navigate to the page.
2. From the *Add* menu (), open *Widgets* → *Collaboration*.
3. Drag a *Message Boards* widget onto the page.
4. Click the *Options* icon ( ) in the widget's title bar and select *Configuration*.

5. From the *Scope* menu in the *Scope* tab, select the page's name or *PageName (Create New)* if the page scope doesn't exist yet.

6. Click *Save*, and then close the dialog.

Note that you must still use the Site Administration menu to administer a page-scoped Message Boards widget. You do so by setting the Site Administration menu's active scope. Follow these steps to do this:

1. Open the *Menu* (), expand the menu for your Site, then expand *Content*.

2. The current scope appears just below the *Content* heading. *Default Scope* is the current Site. To change this, click the gear icon () and then select your desired scope. This changes the Site Administration menu to reflect scope you selected. To work in a page's scope, for example, select that page from the gear icon. That page's name then becomes the Site Administration menu's title.

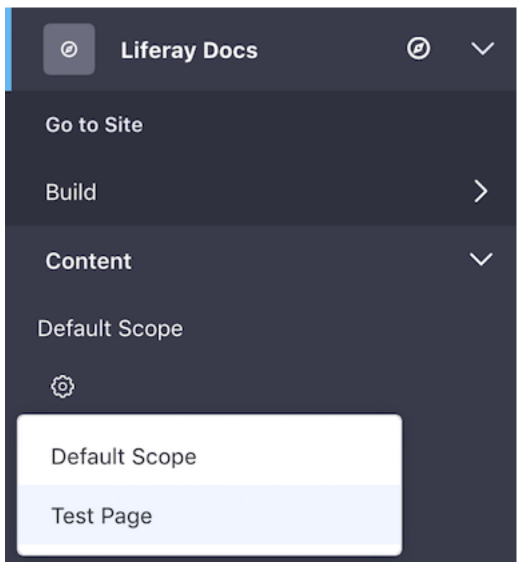

Figure 52.2: Select the page's scope under the *Content* menu in Site Administration.

3. Select *Message Boards* from the *Content* menu. Any changes you make here apply to the scope that you selected in step 2.

Globally-scoped Message Boards

To manage a message board in the global scope, follow these steps:

1. Open the *Menu* (), then click the compass icon () on the Site Administration menu. This opens the Select Site dialog.

2. Select the *My Sites* tab, then select *Global*. This closes the dialog and changes the Site Administration menu's title to *Global*.

3. Select *Message Boards* from the *Content* menu. Any changes you make here apply to the global scope.

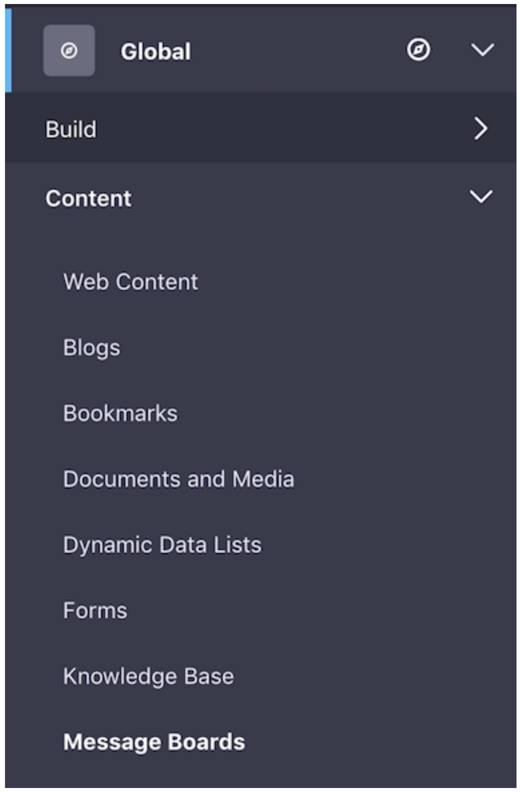

Figure 52.3: After changing to the global scope, select *Message Boards* from the *Content* menu in Site Administration.

52.2 Configuring Message Boards

Before using a message board, configure it to your needs. First, open the Message Boards app in your scope's Site Administration menu, as described earlier. To open the message board's configuration screen, click the message board's *Options* menu () and select *Configuration*. The below sections cover these tabs.

General Setup

The *General* tab contains general settings:

Allow Anonymous Posting: Choose if users can post anonymously. Use this with caution—anonymous users tend to be mean.

Subscribe by Default: Choose if users are subscribed automatically to threads in which they've posted.

Message Format: Define the markup language of users' message board posts. You can choose BBCode or HTML. When creating posts, the type of WYSIWYG editor presented to users depends on which option is enabled. Both editors have a *Source* button that lets users view a message's underlying BBCode or HTML. Users can compose messages using either the WYSIWYG or Source view and can switch between views during message composition by clicking the *Source* button. For security reasons, BBCode is preferred.

Enable Report Inappropriate Content: Choose if users can report content as inappropriate. This sends a message to administrators so they can take action.

Enable Ratings: Choose if users can rate posts.

Thread as Question by Default: This automatically checks the *Mark as question* box in the new thread window. Threads marked as questions display *waiting for an answer*. Replies to the original message can be marked as an answer.

Show Recent Posts from Last: The *Recent Posts* tab shows posts from the following timeframes you define here:

- 24 hours
- 7 days (default)
- 30 days
- 365 days

After the allotted time has passed, the post expires from *Recent Posts*, but is still accessible everywhere else in the message board.

Email Setup

Use these tabs to configure how the Message Boards app handles email notifications:

Email From: The name and email address that sends email notifications. The default administrator account's name and email address are from the `admin.email.from.name` and `admin.email.from.address` portal properties. These were set in the Basic Configuration Wizard when installing Liferay DXP. Make sure to update this email address to a valid one that can be dedicated to notifications.

HTML Format: Support HTML in these emails.

Definition of Terms: Shows variables you can use in the email templates you'll define next.

Message Added Email: Create a template for email users receive when a message is added to a topic they subscribe to.

- **Enabled:** Whether automatic emails are sent to subscribed users.
- **Subject:** Choose a prefix for the email's subject line. This lets users set up message filters in their email clients for these notifications.
- **Body:** The message body content. Use the variables defined in *Definition of Terms* to customize this content for users.
- **Definition of Terms:** Shows variables you can use in the email templates.

Message Updated Email: This tab is identical to the Message Added Email tab, except it defines the email that users receive when a post is updated.

Thread Priorities

The *Thread Priorities* tab defines custom priorities for message threads. This lets privileged Roles tag a thread with a certain priority, which highlights it for users. Three priorities are defined by default:

- Urgent
- Sticky
- Announcement

To define a thread priority, enter its name, a URL to its image icon, and a priority number. Threads with a higher priority are posted above threads with a lower priority.

Thread Icons

Icon	Definition
	Urgent
	Announcement
	Sticky
	Question

The localized language field lets you name the priorities in each locale. You can select the locale, update the priority names for it, and save your updates.

User Ranks

The User Ranks tab ranks users by the number of messages they have posted. Default ranks from 0 to 1000 are provided, but you can set custom ranks here as well.

You can also use this to define message boards labels that appear on user profiles. For example, you can use the message boards label *Moderator* for anyone who is a part of any of the Message Boards Administrator groups: the Site Role, the Organization, the Organization Role, the regular Role, or the User Group:

```
Moderator=organization:Message Boards Administrator

Moderator=organization-role:Message Boards Administrator

Moderator=regular-role:Message Boards Administrator

Moderator=site-role:Message Boards Administrator

Moderator=user-group:Message Boards Administrator
```

As with thread priority names, the *Localized Language* field localizes rank names.

RSS

Message board threads can be published as RSS feeds. The RSS tab enables/disables RSS subscriptions and defines how the feeds are generated:

Maximum Items to Display: The number of items to display in the feed.

Display Style: The feed's appearance. You can publish the full content, an abstract, or just the thread title.

Format: The feed's format: RSS 1.0, RSS 2.0, or Atom 1.0.

Once you've finished configuring your message board, make sure to *Save* your changes.

52.3 Message Board Permissions

Open the Message Boards app in your scope's Site Administration menu, as described in the article on creating message boards. Then click the *Options* icon () and select the *Home Category*

Permissions option. This permissions screen is for granting and revoking access to message board functions.

The permissions enable a Role to perform the following actions:

Lock Thread: Stop any further additions or modifications to a thread's messages.

Subscribe: Receive notifications on new and modified posts.

Reply to Message: Respond to an existing message.

Add File: Attach a file to a message.

Permissions: View and modify permissions.

Add Message: Post a new thread.

View: View all the contents of message threads.

Add Category: Add a new category to the message board.

Update Thread Priority: Modify a thread's priority.

Ban User: Forbid a user from participating in the message board.

Move Thread: Move a thread to a different category or subcategory.

Configure the Roles with the permissions you want and *Save* your changes.

After adding a Message Boards widget to a page, you can access that widget instance's general permissions. To do so, select the widget's *Options* menu ( ) and select *Permissions*. This permissions screen lets you control access to the widget instance's Permissions, Preferences, and Configuration menus.

52.4 Message Board Categories

Message Board categories organize threads by topic. This makes it easier to find the right topic for discussion, and can also help discussions stay on topic. For example, a tropical fishkeeping message board may have separate categories for freshwater and saltwater topics.

This article shows you how to create and manage message board categories.

Adding Categories

Follow these steps to create a message board category:

1. Open the Message Boards app in your scope's Site Administration menu, as described in the article on creating message boards.

2. Click the *Add* icon ( ) and select *Category*. This opens the Add Category form.

3. Enter a name and description for the category.

4. Select the category's *Display Style*. This controls how threads in the category appear. By default, you can choose these display styles:

    - **Default:** Classic display style for general purpose discussions.
    - **Question:** Threads appear in a question and answer style.

    You can create custom display styles and make them available for selection in this form. You must set the available display styles via the portal property `message.boards.category.display.styles`. Similarly, you can set the default display style in `message.boards.category.display.styles.default`.

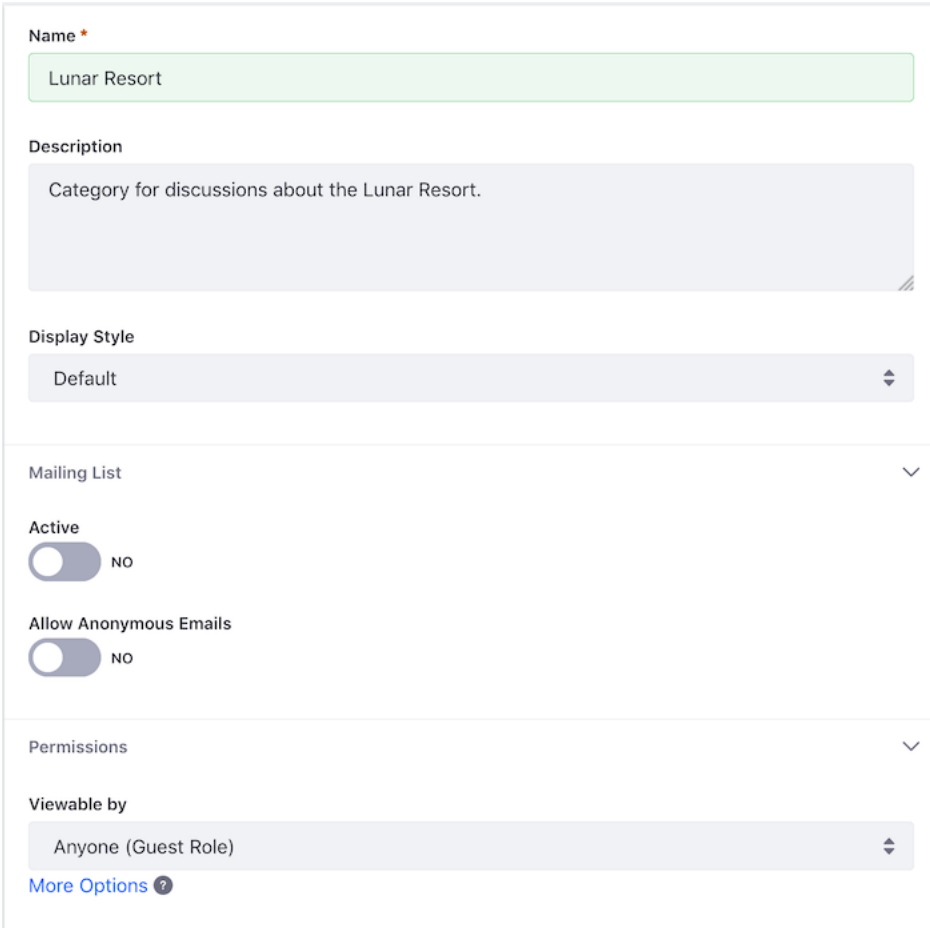

Figure 52.4: You have several options to create a message board category for your needs.

5. Open the *Mailing List* section of the form and set the mailing list options you want. To enable a mailing list for the category, set the *Active* toggle to *YES*. To enable anonymous emails in the list, set the *Allow Anonymous Emails* toggle to *YES*. The default for both toggles is *NO*. For an explanation of these features, see the documentation on mailing lists for Message Boards.

6. Open the *Permissions* section and set the category's permissions. The *Viewable by* selector lets you pick who can view the category:

    - Anyone (Guest Role)
    - Site Members
    - Owner

    To show more permissions options, click *More Options*. A table appears with the rest of the category's permissions, which you can assign to the Guest and Site Member roles:

    Lock Thread: Stop any further additions or modifications to a thread's messages.

    Add Subcategory: Add a new category within this category.

    Update: Edit the category.

    Subscribe: Receive notifications on new and modified posts.

Reply to Message: Respond to existing messages.

Add File: Attach a file to any of your messages.

Permissions: View and modify permissions.

Delete: Remove the category.

Add Message: Post a new thread.

Update Thread Priority: Modify a thread's priority.

Move Thread: Move a thread to a different category or subcategory.

Note that after creating a category, you can revisit its permission options by clicking the category's *Actions* icon (⋮) and selecting *Permissions*.

7. Click *Save* when you're finished. Your category now appears in the table.

As you add categories to a message board, they appear on the message board's home screen. The list displays the category names and the numbers of subcategories, threads, and posts in each one.

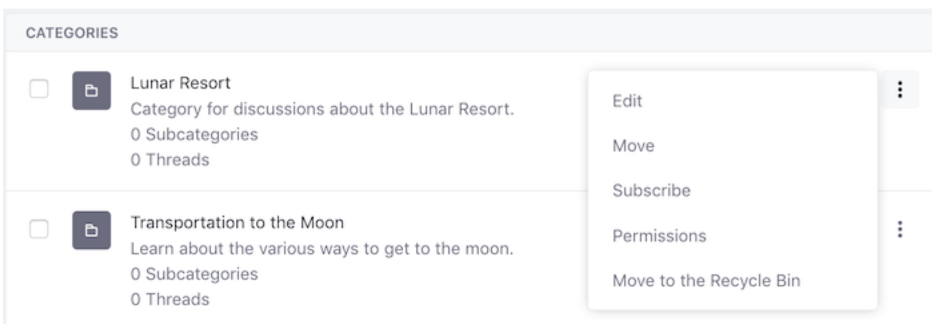

Figure 52.5: Categories help you organize threads so users can find topical threads that interest them.

Adding Subcategories

Categories can contain as many subcategories as you like. If, however, you nest categories too deep, users can have trouble finding them.

Follow these steps to add a subcategory to a category:

1. Click the category's name in the list, then click the *Add* icon (+) and select *Subcategory*. The Add Subcategory form appears.

2. Fill out the Add Subcategory form with the values and settings you want to use for the subcategory. This form is identical to the Add Category form, and is populated with the parent category's properties by default.

3. Click *Save* when you're finished. Your subcategory now appears in the table.

Moving and Merging Categories

Each category can have any number of threads, and you can add as many categories and subcategories as you wish. You can also move and merge categories.

Follow these steps to move a category or merge it with another:

1. Click the category's *Actions* icon (⋮) and select *Move*. This brings up the Move Category form.
2. Select a new parent category via the *Select* button under the *Parent Category* field. Note that this field is empty for top-level categories.
3. If you want to merge the category with the selected parent category, select *Merge with Parent Category*.
4. Click *Move*.

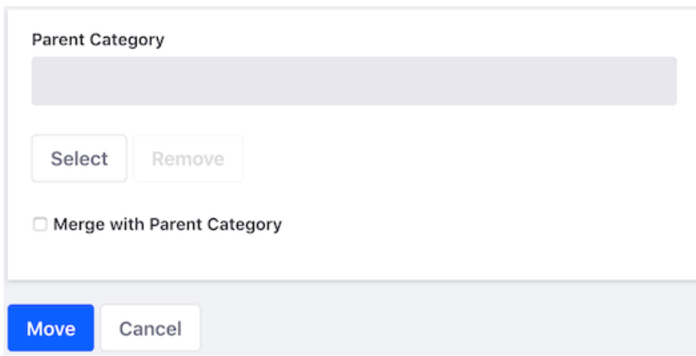

Figure 52.6: The Move Category form lets you move and merge categories.

52.5 User Subscriptions and Mailing Lists

The Message Boards notifies users of message boards activity via email in two ways:

- User subscriptions
- Mailing lists

Note: Since multiple sites can use a globally scoped message board, such message boards don't support user subscriptions or mailing lists. Make sure to use a site-scoped or page-scoped message board if you need user subscriptions or a mailing list.

User Subscriptions

In the user subscriptions mechanism, the Message Boards app uses its configured *Email From* address to send email notifications to subscribed users. The app can also import email replies to message board notifications directly into the message board. Then, users can interact on the message board via email without logging in to view the message board directly. This is disabled by default. To enable it, add the following line to your `portal-ext.properties` file:

```
pop.server.notifications.enabled=true
```

The user subscription mechanism uses the POP mail protocol. When the Message Boards app receives an email reply to a message board notification, it posts that reply to the message board and then deletes it from the mail server. Deleting the message from the mail server is the POP protocol's default behavior and the Message Boards app assumes that your POP mail server behaves this way. Most POP clients offer an option to leave mail on the mail server after it downloads, but you shouldn't exercise this option. If you configure mail to be left on the mail server, the Message Boards app sends copies of each retained message along with each new email notification it sends to subscribed users.

When enabling Message Boards to import replies to email notifications, you must decide whether to handle notifications with a mail server subdomain. By default, the following property setting is specified in the portal properties:

```
pop.server.subdomain=events
```

This property creates a special MX (mail exchange) subdomain to receive all virtual instance related email (e.g., events.liferay.com). If you don't want to use this approach, unset this value in a `portal-ext.properties` file:

```
pop.server.subdomain=
```

Doing so tells Message Boards to use the *Email From* address specified in the Message Board's configuration to receive message board notification email replies. For example, the *Email From* address could be set to *replies@example.com*.

If you're not using a mail subdomain, Message Boards parses the message headers of emails from the *Email From* address to determine the message board category and message ID. If you keep the `pop.server.subdomain=events` default, the email notification address takes the following form:

```
mb.[category_id][message_id]@events.liferay.com
```

In this case, Message Boards parses the email address to find the category and message ID. Parsing the email address is safer than parsing message headers, since different email clients treat message headers differently. This is why the `events` subdomain is enabled by default.

You can also configure the interval on which the `POPNotificationListener` runs. The value is set in one minute increments. The default setting is to check for new mail every minute, but you can set it to whatever you like:

```
pop.server.notifications.interval=1
```

Note: Depending on your mail provider, if you use multiple devices to access email through POP, you might need to configure in your POP settings something like Gmail's *recent mode*. In Gmail, recent mode assures that emails go to all your devices instead of only the first client that receives the email. To enable recent mode in Gmail, for example, prefix the value of your POP client's Username or Email field with `recent:`.

Mailing Lists

Alternatively, the Message Boards app can use mailing lists to send email notifications. Any category in a message board can have its own mailing list. The mailing list mechanism, unlike the user subscription mechanism, supports both the POP and the IMAP protocols. POP is the default, but each message board's mailing list is configured independently. If you choose the IMAP protocol for a category's mailing list, make sure to configure the IMAP inbox to delete messages as they are pulled by the email client that sends messages to the users on the mailing list. Otherwise, each email message retained on the server is sent to the mailing list each time there's a new post or update in the category.

When a mailing list is enabled for a message board category, Message Boards listens to the specific email inbox that's configured for the mailing list. Enabling the mailing list function lets users on the mailing list reply to the notification messages in their email clients. Message Boards pulls the messages from the email inbox it's configured to listen to and automatically copies those replies to the appropriate message board thread.

To enable the mailing list functionality for a category, follow these steps:

1. Set up a dedicated email address for the category.

2. Click the category's *Actions* icon ( ) and select *Edit*.

3. In the *Mailing List* section of the form, set the *Active* slider to *YES*. Several options then appear. Fill these out as follows:

   Email Address: The email address of the account that receives the messages.

   Protocol: Select POP or IMAP.

   Server Name: Your mail server's host name.

   Server Port: The port on which your mail service is running.

   Use a Secure Network Connection: Whether to use an encrypted connection if your server supports it.

   User Name: The login name on the mail server.

   Password: The password for the account on the server.

   Read Interval (Minutes): How often to poll the server looking for new messages to post.

   Email Address (Outgoing): The email address originating messages from this category. If you want your users to be able to reply to the categories using email, this should be the same address as the incoming email address.

   Use Custom Outgoing Server: Use a different mail server than global default. Fields appear for configuring the server's name, port, user name, password, and secure connection.

4. If you want to let emails from anonymous users post to the message board category, set the *Allow Anonymous Emails* toggle to *YES*.

5. Click *Save* when you're finished.

52.6 Using the Message Boards

You can add a Message Boards widget to a page from the *Add* ( ) menu's *Collaboration* section. The Message Boards interface is similar to other message boards that populate the Internet. In any case, it can't hurt to explore how to use Liferay DXP's Message Boards and discover its features.

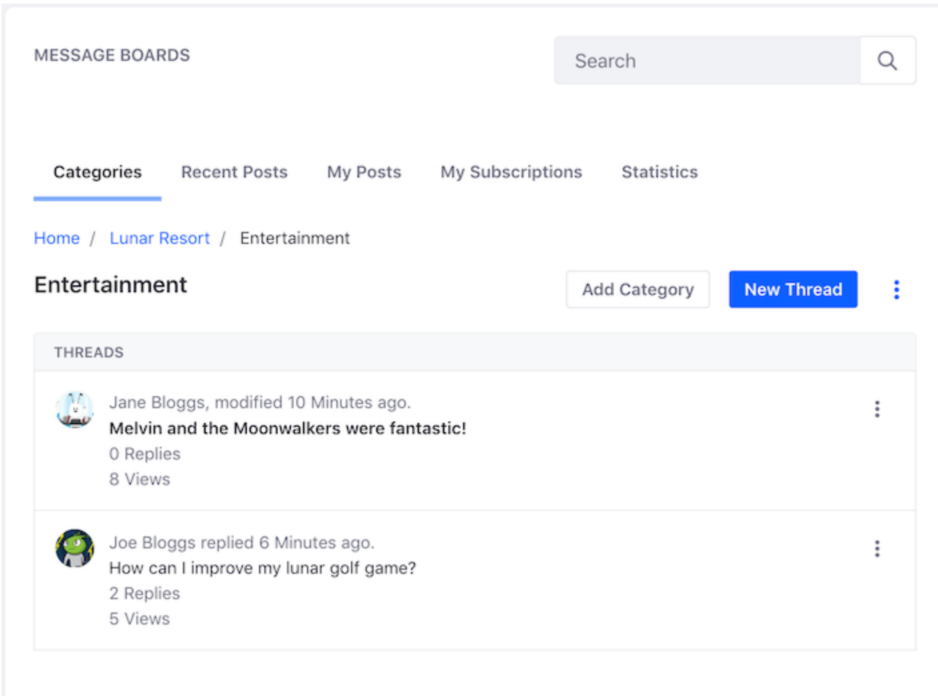

Figure 52.7: The Message Boards widget lets you explore its categories, interact with message threads, and post new messages.

Threads can be viewed many ways. At the top of the widget is a set of tabs:

- **Categories:** The message board's categories.
- **Recent Posts:** Posts from all categories, sorted by date.
- **My Posts:** The current user's posts.
- **My Subscriptions:** Lets users view and manage their thread subscriptions.
- **Statistics:** The number of categories, posts, participants, and a list of the top contributors. You can also access this from the same tab in Site Administration's Message Board app.

You can also use the search bar at the top of the widget to search for threads and posts. Although search works on threads and posts within categories, it doesn't work on categories themselves.

Posting New Threads

Follow these steps to post a new thread:

1. Click the Message Boards widget's *New Thread* button. Alternatively, click the *Add* button ( ) and select *Thread* in the Message Boards app in Site Administration. Either way, the same *Add Message* form appears.

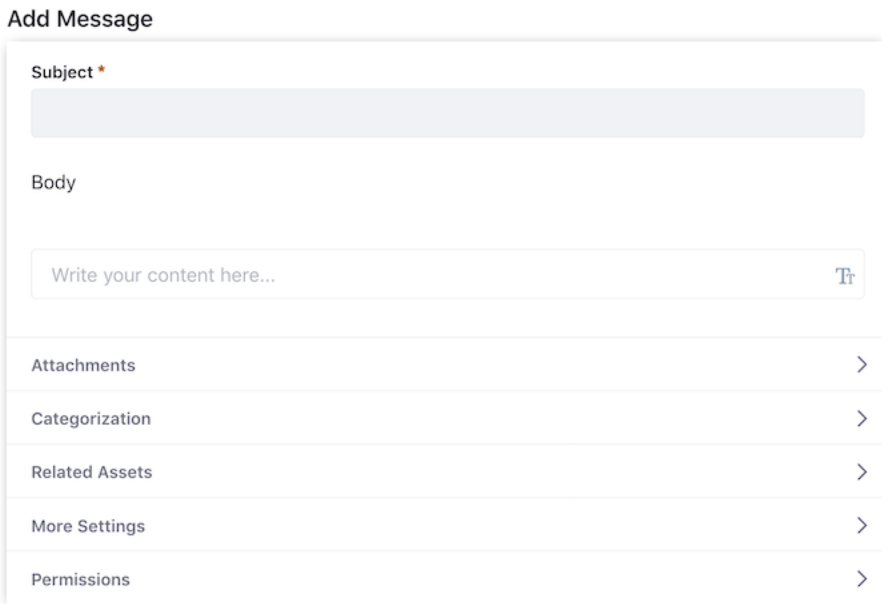

Figure 52.8: The Add Message form lets you create a new thread.

2. Give your thread a title in the *Subject* field.
3. Create your thread's content in the *Body* field. This field uses the same editor as the Blogs app, except that it uses BBCode instead of HTML. For further instructions, see the documentation on using the editor. Also note that you can mention other users by entering the `@` character and their user name.
4. If you want to add attachments, open the *Attachments* section and add them via drag and drop or the *Select Files* button.
5. If you want to associate a tag with the message, open the *Categorization* section and use the *Select* button to select an existing tag. You can also create a new tag by entering it in the *Tags* field and clicking *Add*. See the documentation on tags for more information.
6. If you want to select an existing asset in the portal (e.g., a media file, blog post, etc.) to relate to your thread, open the *Related Assets* section and use the *Select* button to select that asset.
7. Open the *More Settings* section and select the settings you want to use:
   - **Mark as a Question:** Whether to mark this thread as a question. This lets you later select a post in the thread as the answer.
   - **Anonymous:** Whether this thread is posted anonymously.
   - **Subscribe Me:** Receive notifications for activity on the thread.
   - **Priority:** The thread's priority in the Message Board. By default, you can choose *Urgent*, *Sticky*, or *Announcement*. Additional priorities can also be configured in the Message Boards app in Site Administration.
   - **Allow Pingbacks:** Whether pingbacks are allowed for your thread.
8. Open the *Permissions* section and set the thread's permissions. Possible values in the *Viewable by* selector are

- Anyone (Guest Role)
- Site Members
- Owner

You can also click the *More Options* link to select additional permissions:

- **Update:** Edit the thread.
- **Subscribe:** Receive notifications for thread activity.
- **Permissions:** Grant/revoke thread permissions.
- **Delete:** Remove the thread.
- **View:** View the thread.

Note that you can revisit the thread's permissions after posting it. To do so, select the thread's Actions menu (⋮) and select *Permissions*.

9. Click *Publish* to publish the thread. Once it's published, it appears along with the other threads in the category.

Participating in Message Board Threads

To find message board threads that interest you, browse a message board's Categories or Recent Posts tabs. In the Categories tab, you can view a category's thread listing by clicking the category's name. Within a category, you can subscribe to an RSS feed and/or emails that inform you about activity in that category. The Recent Posts tab also lists threads, except they're the latest threads across all categories.

Click a thread to view it. Messages appear in a threaded view so that replies are aligned under their parent thread. This makes it easy to follow conversations. Thread replies are indented under their parent thread.

Subscribing to a thread causes Message Boards to send the user an email whenever a new message is posted to the thread. If you have enabled the mailing list feature for the thread's category, users can reply to these messages to post back to the thread without having to visit your site.

Most threads get more interesting as users reply to them. Follow these steps to reply to a message in a thread:

1. Click the *Reply* button. This opens the quick reply form, which only contains a text field for entering your reply.

2. Enter your reply in the text field. To access more options for your reply, click the *Advanced Reply* link. This opens the full editor from the add/edit thread form.

3. Click *Publish* to publish your reply.

In addition to replying to a message, you can rate it or flag it as objectionable. A message board moderator can evaluate flagged messages and decide how to handle the messages and their authors.

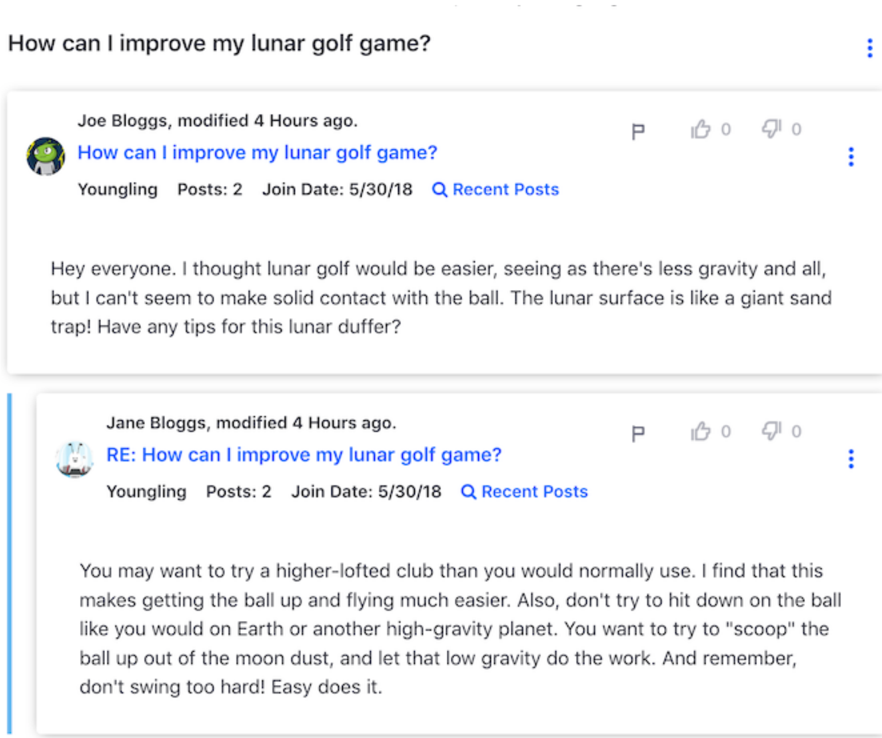

Figure 52.9: A thread's view displays author information and thread content, for the thread and all replies to the thread.

52.7 Managing Message Boards

Message boards can become unwieldy if left unmanaged. The Message Boards in Site Administration facilitates day-to-day thread administration. You may wish to assign this function to a Role that you give to one or more users. This frees you to concentrate on other areas of your Site. For example, you can create a Role called *Message Board Administrator* scoped to the portal (globally), an Organization, or a Site. Members of a global Role can administer Message Boards throughout the portal. Members of an Organization or Site-scoped Role can only administer Message Boards in that Organization or Site, respectively.

Follow these steps to create a global Role:

1. In the Control Panel, select *Users* → *Roles*.

2. Select or create the Role.

3. Select the Role's *Define Permissions* tab and navigate to *Site Administration* → *Content* → *Message Boards*. A screen appears that lets you configure Message Board permissions.

4. Select the permissions you want message board administrators to have, then click *Save*.

5. Add users to this Role.

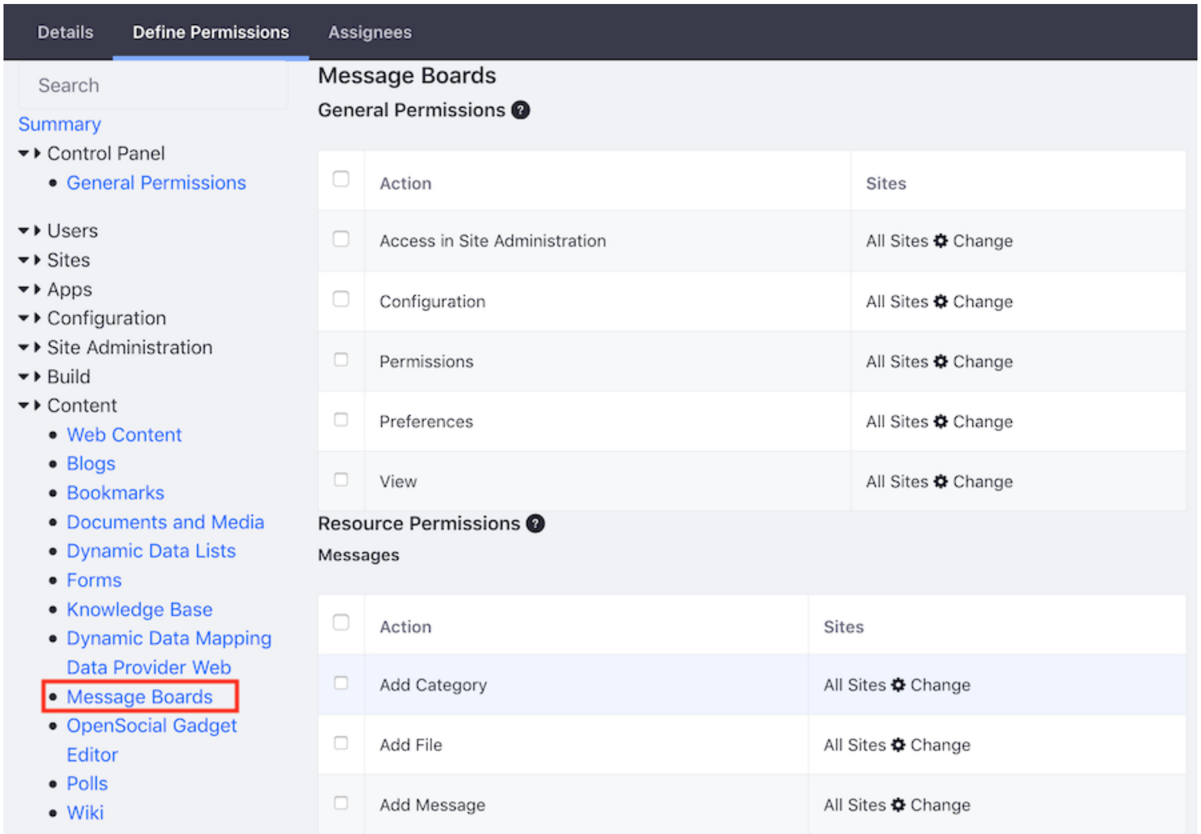

Figure 52.10: Define the permissions you want to use for the message boards administrators.

Locking Threads

You may encounter threads that you think should be preserved, but stopped. You can halt activity on a thread by selecting *Lock* from the thread's *Actions* menu (⋮).

Moving Threads

If someone posts a thread to the wrong category, you can move it to the proper one. Follow these steps:

1. Select *Move* from the thread's *Actions* menu (⋮). This opens the *Move Thread* form.
2. Click the *Select* button and select the new category.
3. If you want to add a post explaining the move, select *Add explanation post*.
4. Click *Move* to move the thread.

Deleting Threads

Sometimes users begin discussing topics that are inappropriate or that reveal confidential information. In this case, administrators can delete the thread from the message boards. To do so, select *Move to Recycle Bin* from the thread's *Actions* menu (⋮).

Banning Users

Unfortunately, message board users can be abusive. In this case, you can ban the user from the message boards. While viewing any of the user's posts in any thread, select the post's *Actions* menu (⋮) and select *Ban this User*.

To reinstate a banned user, you must use the Message Boards app in Site Administration. Navigate to this app and select the *Banned Users* tab. Select the user's *Actions* menu (⋮) and select *Unban this User*.

Splitting Threads

Sometimes a thread goes on for a while and the discussion completely changes into something else. In this case, you can split the thread where the discussion diverged and create a whole new thread for the new topic. To split a thread at a certain post, administrators can select that post's *Actions* menu (⋮) and select *Split Thread*. This brings up a form that lets you add an explanation post to the split thread. Click *OK* to split the thread.

Editing Posts

Administrative users can edit anyone's posts, not just their own. Sometimes users post links to copyrighted material or unsuitable pictures. By editing these posts, you can redact information that shouldn't be posted, or remove content not conforming to your terms of use. You can also update the thread's priority or mark a reply as an answer to a thread's question.

To edit a post, select its *Actions* menu (⋮) and select *Edit*.

Post Permissions

Permissions can be set not only on threads, but also on individual posts. You can choose to limit a particular conversation or post to only a select group of users. To do this, select the post's *Actions* menu (⋮) and select *Permissions*. You can then choose which Roles have the following permissions:

- Delete
- Permissions
- Subscribe
- Update
- View

Use this, for example, to let some privileged users post on a certain thread, while others are only allowed to view it. Other combinations of these permissions are also possible.

52.8 Mentioning Users

Have you ever wanted to include another user in a discussion on the Message Boards? Have you ever wanted to give kudos to a colleague in content you're writing? You can mention (notify and/or draw attention to) other users by entering the `@` character in front of each user's user name.

When you mention a user, the user receives a site notification next to the user's profile icon and an email, alerting the user with a link to the content. You can mention users in a blog entry, a message boards post, or comments in any app that supports comments. A mention also links to the user's home page, so readers can find out more about that user.

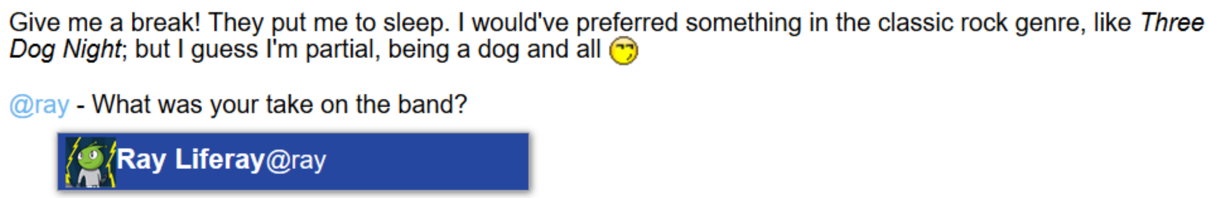

Figure 52.11: As you enter a user name after the `@` character, links to users that match the text you enter are displayed. Select the user you want to mention and publish your content.

A selector appears after entering the `@` character, listing users that match the name you're entering. In the selector, users are represented by their profile picture, name, and user name. Click the user you want to mention and finish editing your content.

On publishing the content, mentioned users receive a notification next to their profile picture and an email, informing them that they've been mentioned. The notification and email indicate the author's name and content type, and it links to the content.

Figure 52.12: The number of notifications a user has appear in a badge on the user's profile.

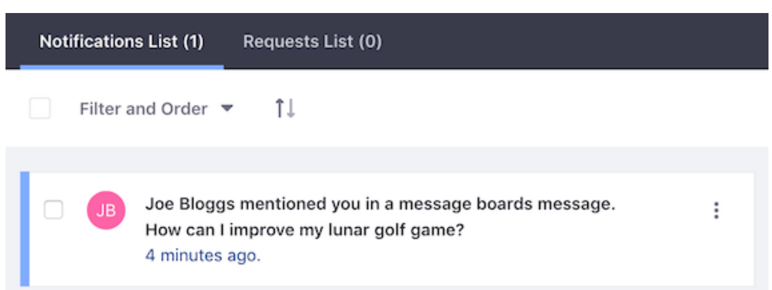

Figure 52.13: When you click the notifications number next to your profile picture, your Notifications List appears.

Mentions are enabled by the Mentions app, which is a part of the Collaboration Suite. By default, the Mentions app is enabled globally. However, you can enable/disable it globally or per site. For a site to use Mentions, it must be enabled for the site's Virtual Instance.

To access the global Mentions settings for your Virtual Instance, open the *Menu* (), navigate to *Control Panel* → *Configuration* → *Instance Settings*, click the *Social* tab, and expand the *Mentions* section.

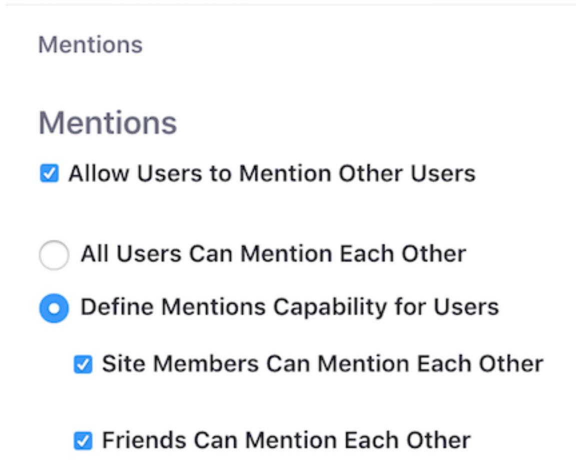

Figure 52.14: From Instance Settings in the Control Panel, you can enable or disable the Mentions feature for all of the Virtual Instance's sites.

By default, all users are allowed to mention fellow site members and friends. To fine tune these options, select *Define Mentions Capability for Users* and specify the settings you want.

For Mentions to be available for a site, the app must be enabled for a site's Virtual Instance. Site administrators can enable or disable Mentions for a site. A site's Mentions app configuration is accessible from within the *Menu* (). Once in the menu, navigate to *Site Name* → *Configuration* → *Site Settings*. In the *Social* tab, expand the *Mentions* section. Enable or disable mentions via the toggle labeled *Allow Users to Mention Other Users*.

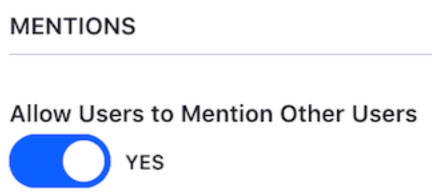

Figure 52.15: Mentions can also be enabled or disabled per site.

Chapter 53

Working Together with the Wiki

Wikis are for collaboratively building a collection of information. The most famous wiki on the planet is Wikipedia. It's a full encyclopedia developed collaboratively by users from all over the world, using a wiki. Liferay DXP's wiki does these things:

- Creates multiple wikis in a single wiki app instance.
- Scopes wikis to a page, a site, or the entire portal.
- Creates and edit wikis in WikiCreole syntax.
- Attaches files to wiki articles.
- Associates wiki articles with other assets in the portal.
- And more.

As you can see, Liferay DXP's wiki is flexible and can be configured to fit nearly any use case. What's more, it's completely integrated with the portal's user management, tagging, and security features.

53.1 Getting Started with Wikis

The Menu () is the best place to start working with your wikis. Click the *Menu* (), navigate to your site, and select the *Content* section. If you're updating an existing page-scoped wiki app instance, you can select that page scope from the scope menu the Gear icon () makes available. The site's wiki app instance is available in the Default scope. Once you're in the proper content scope, click *Wiki*. The Wiki administration screen lets you add, modify, and delete wiki nodes. A Wiki app instance can contain many wiki nodes. By default, it contains one node: *Main*.

Configuring Wikis

Before adding to your wiki instance, you should configure it. The instance's interfaces for permissions, export and import, configuration, and application templates are accessible from the Options menu. Click the *Options* icon () to open this menu.

The following options are available in this menu:

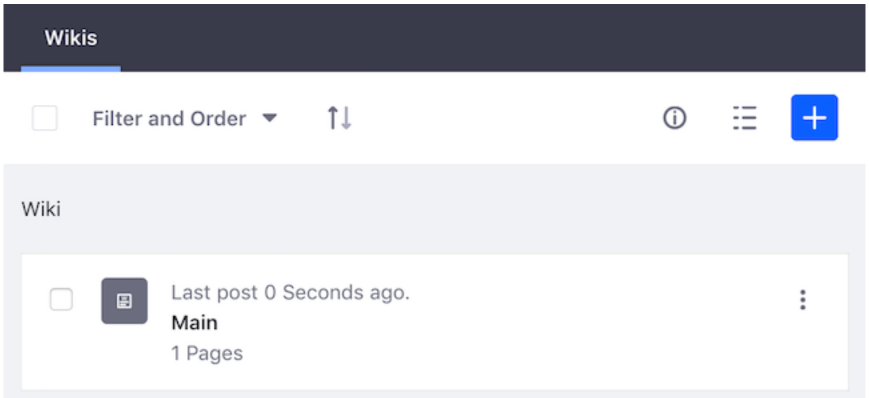

Figure 53.1: The Wiki app instance has a wiki node named *Main* with a single front page. You can build on the Main node or click the Add icon to create a new node.

Wikis Permissions: Specify which Roles can create wiki nodes and access the Wikis Permissions screen. For example, if you've created a specific Role for creating wiki nodes and want to enable that Role to create new wiki nodes in this wiki application instance, select the Role's check box in the *Add Node* column and then click *Save*.

Export / Import: Import existing wiki content into your wiki app instance, or export wiki content to a file. For details, refer to Importing/Exporting Pages and Content.

Configuration: Configure email notifications and RSS feeds. The *Email From*, *Page Added Email*, and *Page Updated Email* tabs are similar to other apps' notification email settings tabs; they customize who wiki emails come from and the format and text of the email sent when a page is added or updated. The *RSS* tab lets you configure RSS feeds.

Permissions: Specify which Roles can view the wiki app instance's Options menu, access the menu's Configuration and Permissions options, and access any custom preference options added to the wiki app.

Adding Wikis

Follow these steps to create a new wiki node:

1. Click the *Add* icon (+) to start creating a new wiki node. The *New Wiki Node* form appears.

2. Add a name and description for the wiki node.

3. Open the form's *Permissions* section and define the wiki node's permissions. You can select the following permissions in the *Viewable by* menu:

   - Anyone (Guest Role)
   - Site Members
   - Owner

   You can also click the *More Options* link to assign permissions to specific Roles.

4. Click *Save* when you're done creating the wiki node.

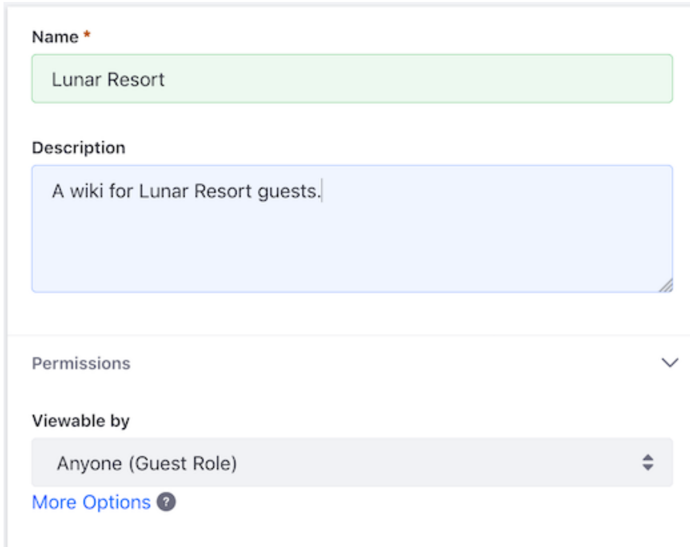

Figure 53.2: The New Wiki Node form lets you describe your new node, set view permissions, and set permissions for the Guest and Site Member roles.

Wiki Node Options

Next to each listed wiki node is an *Actions* menu ( ). Here are the actions available in this menu:

Edit: Edit the wiki's name and description.

Permissions: Specify which roles can add attachments to wiki pages, add pages, delete pages, import pages, set permissions on the wiki node, subscribe to modifications, update existing pages, and view the wiki node.

Import Pages: Import data from other wikis. This lets you migrate from another wiki application to the Liferay DXP wiki. You might want to do this if you're migrating your site from a set of disparate applications (i.e., a separate forum, a separate wiki, a separate content management system) to Liferay DXP, which provides all of these features. Currently, MediaWiki is the only supported wiki.

RSS: Subscribe to an RSS feed using Live Bookmarks, Yahoo, Microsoft Outlook, or an application on your machine.

Subscribe: Subscribe to a wiki node. Any time a wiki page is added or updated, the portal sends you an email notification.

View Removed Attachments: Display attachments that have been removed from the wiki node.

Move to Recycle Bin: Moves the wiki node to the Recycle Bin.

Before opening wiki nodes to contributors, you should consider whether to associate a workflow with them. For example, you could create a workflow that requires an administrator's approval to publish a wiki page modification (add, update, or delete). You can access your site's default *Wiki Page* workflow from within the Site Administration Menu, by navigating to *Configuration* → *Workflow* for your site. To learn how to use workflow, see the Workflow section.

53.2 Adding and Editing Wiki Pages

Wiki nodes initially have no pages. When you navigate into a node for the first time, a default page called *FrontPage* is created automatically. To view the page, click the wiki node's name and then

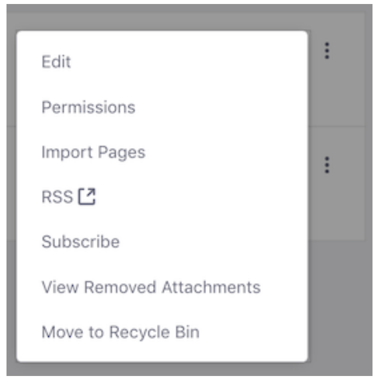

Figure 53.3: Each wiki node's Actions menu lists actions you can perform.

click *FrontPage*. The FrontPage appears and shows a message that explains the page is empty and needs you to add content. That message is a link; click it to start editing the page. The wiki page editing form then appears.

Note: See the getting started article for instructions on accessing your wiki nodes.

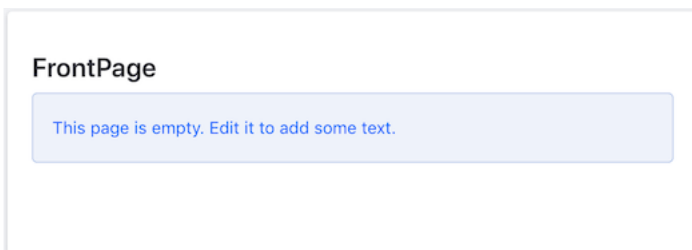

Figure 53.4: Each empty wiki page presents a default message link you can click to edit the page.

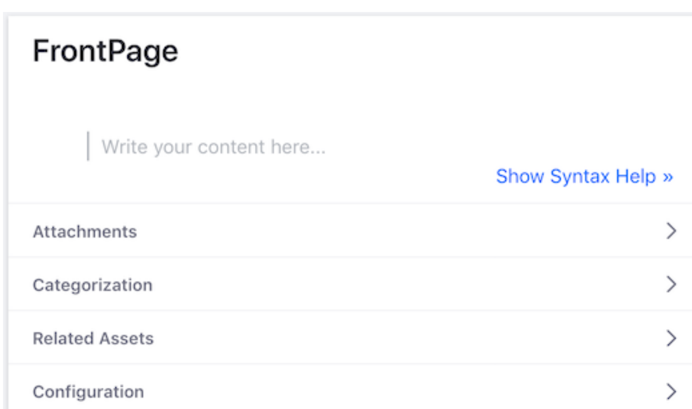

Figure 53.5: The wiki page editing form lets you create and edit your page's content.

Follow these steps to use the wiki page editing form:

1. Enter your content in the field that contains the text *Write your content here....* This is a rich-text, WYSIWYG editor that is almost identical to the one used in the Blogs app. The only difference is that the wiki editor uses Creole instead of HTML as its source. Click the link *Show Syntax Help* if you need help with Creole syntax (e.g., syntax for text styling, header formatting, link creation, etc.). For a detailed explanation of the rest of the editor, see the Blogs documentation.

2. If you want to attach files to the page, open the *Attachments* section of the form and add them via drag and drop or the *Select Files* button.

3. If you want to associate a tag with the page, open the *Categorization* section and use the *Select* button to select an existing tag. You can also create a new tag by entering it in the *Tags* field and clicking *Add*. See the documentation on tags for more information.

4. If you want to select an existing asset in the portal (e.g., a media file, blog post, etc.) to relate to the page, open the *Related Assets* section and use the *Select* button to select that asset.

5. In the form's *Configuration* section, you can set the page to use Creole (default), MediaWiki, plain text, or HTML. We recommend that you stick with the Creole format, as it allows for a much cleaner separation of content and code. You can also use the Configuration section to summarize your edit, and specify whether it's a minor edit.

6. Click *Publish* to publish the page when you're done editing it.

As is common with wikis in general, if you link to a page that doesn't exist, clicking that link opens the new page form with a note stating that the page doesn't exist and that you are creating it.

Note: When you create a page by clicking a link to a page that doesn't exist, the new page is **not** a child of the current page. The page is created at the wiki node's root. From Wiki in Site Administration, you can use the page's Move action to assign it a new parent page. Clicking the Move action brings up a window that lets you select a new parent for the wiki page.

Return to the wiki node view to see a list of the node's top-level pages. If you navigate to a page that has child pages, its child pages are listed. In these page listings, each page's Actions menu ( ) lists the following actions you can take on the page:

Edit: Opens the page in the page editor.

Permissions: Lets you determine which roles can view, update, delete, subscribe to, or set permissions on the page, and add, update, or delete page discussions (comments).

Copy: Opens a page editor window with all the content from the source wiki page. You're prompted to specify a new title for it.

Move: Opens a dialog that lets you rename the page or assign the page to a new parent page within the wiki node.

Subscribe (or Unsubscribe): Subscribes you to (or unsubscribes you from) notifications for the wiki page's modifications.

Move to Recycle Bin: Moves the wiki node to the Recycle Bin.

Each wiki page has a check box next to it. When you select a page's check box, the Management Bar changes to show an Info icon ( ) and Recycle Bin icon ( ). To move the selected page to the Recycle Bin, click the Recycle Bin icon. To get additional information about the page via an info panel, click the Info icon. The info panel provides a star icon that you can select to subscribe to the page's modifications. The info panel's Details section displays the page's summary, format, version, creation and modification dates, number of attachments, and RSS link.

There are several more features in the wiki node view's Management Bar. The *Filter and Order* menu orders the pages by title or modification date and filters them by page type. The arrows button sorts the pages in ascending or descending order. The search bar searches for pages.

The *View Types* button is next to the Info icon. It lets you choose how to display the pages. The View Types button's icon depends on the selected view type:

List ( ): Shows the pages in a list with an icon representing each page. Each page's entry contains the name of its author, when it was last modified, and its workflow status (e.g., Approved, Draft, etc.).

Table ( ): Shows the same information as the List view type, in a smaller list with no page icon. Also, the page's information is in columns and includes the revision number.

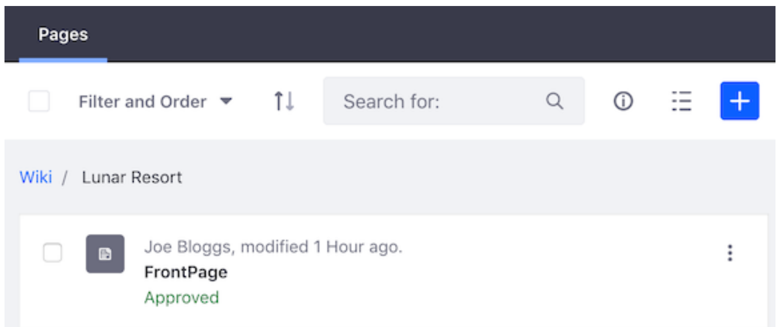

Figure 53.6: The wiki node's view in site administration has features that help you access and learn information about a wiki node's pages.

53.3 Using the Wiki on Site Pages

You can use the Wiki on Site pages via the Wiki widget. Follow these steps to add the Wiki widget to a page:

1. Navigate to the page where you want to place a wiki.

2. From the *Add* ( ) menu, open *Widgets* → *Wiki* and add a *Wiki* to the page.

Your Site's wiki nodes appear in tabs across the top of the widget.

To view the Wiki widget's configuration options, click its *Options* icon ( ) and select *Configuration*. The Configuration screen appears with these tabs:

Setup: Lets you choose wikis to display and gives you several options for displaying them. The *Show Related Assets*, *Enable Page Ratings*, *Enable Comments*, *Enable Ratings for Comments*, and *Enable Highlighting* check boxes enable or disable those features for the Wiki. You can set how you want users to interact with wiki documents. The *Display Template* selector menu lets you choose the Wiki's Application Display Template. Below this, you can set which wiki nodes are visible. For example, you might host two wikis on a given site, exposing one to the public and keeping the other private for site members.

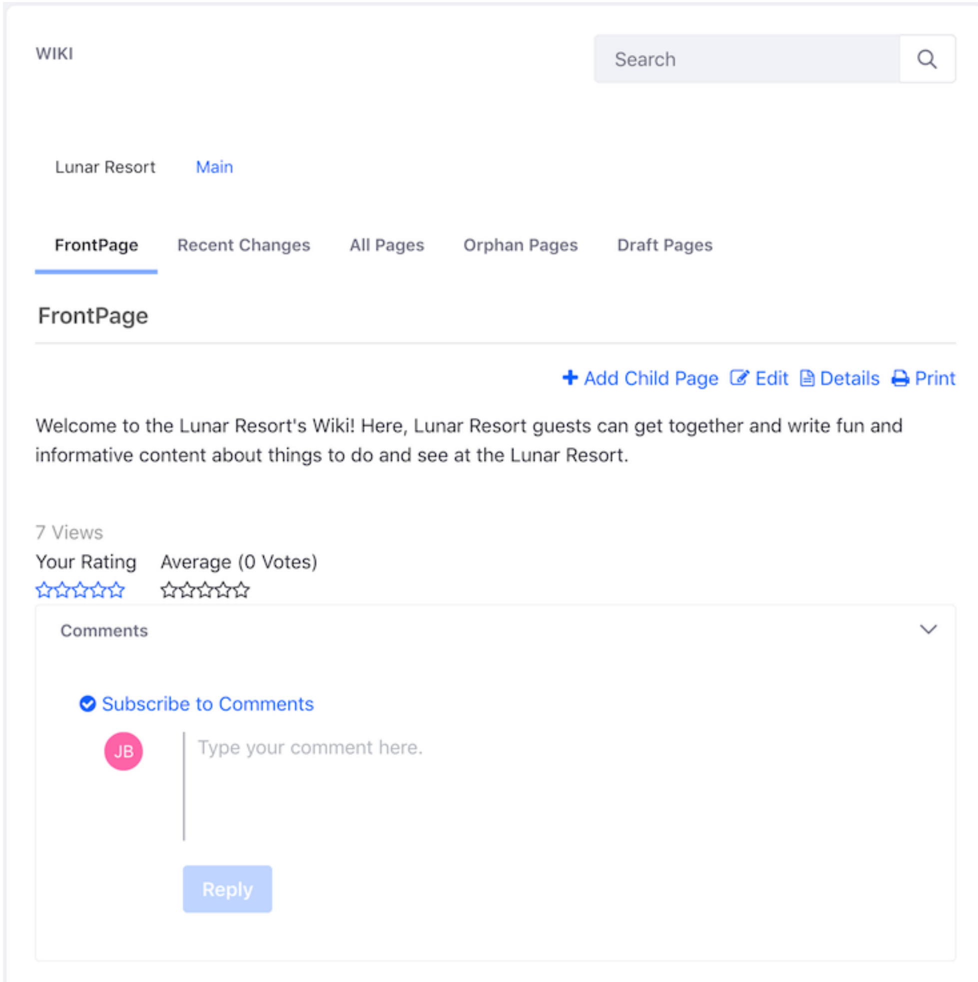

Figure 53.7: Users can interact with your Wiki nodes when you add the Wiki widget to a page.

Communication: Configure communication across portlets, using predefined public render parameters. From here you can modify six public render parameters: `categoryId`, `nodeId`, `nodeName`, `resetCur`, `tag`, and `title`. You can perform these actions on each parameter:

- Ignore the values for this parameter that come from other portlets. For example, the wiki can be used along with the tags navigation app. When a user clicks on a tag in tags navigation, the wiki shows a list of pages with that tag. In some cases, an administrator may want the wiki to show the front page always, independently of any tag navigation done through other portlets. This can be achieved by selecting *Ignore*, so that the values of the parameter coming from those other portlets are ignored.

- Read the value of a parameter from another app. This is an advanced but very powerful option that lets portlets communicate without prior configuration. For example, imagine that the wiki is used to publish information about certain countries, and a custom app that allows browsing countries for administrative reasons was written and placed on the same page. You could associate to this second app a public render parameter called *country* to designate the name of the country. Using this procedure, you can cause the wiki to show the information from the country being browsed in the other app. You can do this here for the wiki by setting the value for the title parameter to be read from the country parameter of the other app.

Sharing: Displays options you're likely to be familiar with such as the sections for sharing the Wiki with websites, Facebook, and NetVibes.

Scope: Set the wiki's scope. You can select the site-scoped or global-scoped instance, or select/create an instance for the page. If the page doesn't already have an instance scoped to it, you can click the *[page name] (Create New)* menu option to create a page-scoped wiki instance.

Once you set the wiki's configuration options the way you want them, click *Save*.

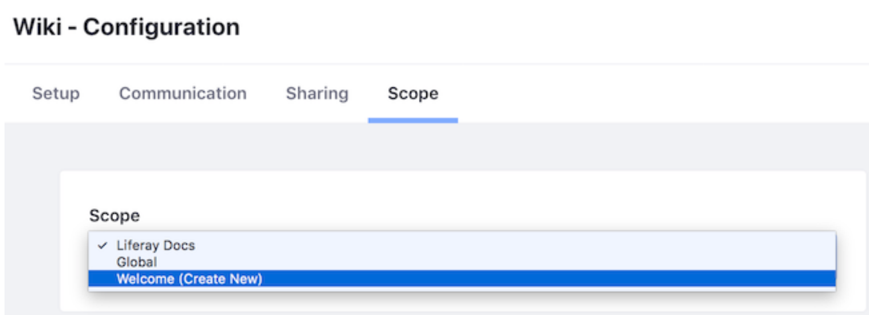

Figure 53.8: Here the user has selected to create a new Wiki instance scoped to the current page named *Welcome*

The Wiki's Options menu also contains the usual widget options:

Look and Feel Configuration: Set the widget's look and feel.

Export/Import: Export or import widget data.

Permissions: Set the widget's permissions.

Configuration Templates: Use configuration templates to store the widget's current setup or apply an existing archived setup.

Remove: Remove the widget from the page.

The Wiki displays links to all of the Wiki instance's nodes, and provides links for navigating around the wiki. Click on a wiki node's name to begin browsing that node's pages. The following navigation links are listed after the wiki nodes:

FrontPage: The wiki node's front page article. This is shown by default when the node is initially selected.

Recent Changes: Shows all of the recently updated pages.

All Pages: A flat, alphabetical list of all pages currently stored in the wiki.

Orphan Pages: A list of pages that have no links to them. This can happen if you remove a page link without realizing it's the only link to that page. This area lets you review such orphaned wiki pages so that you can re-link or delete them.

Draft Pages: A list of unpublished pages. Users can edit pages and save their changes as drafts. They can come back later to finish their changes and publish them.

The current wiki page's content shows in the wiki's main viewing area. Several features display above the wiki page content, depending on which wiki features are enabled and your permissions:

- **Add Child Page:** Add a wiki page as a child of the current wiki page.
- **Edit:** Edit the wiki page (if you have sufficient permissions).
- **Details:** View the wiki page's details (if you have sufficient permissions). This is explained further in the documentation on page details.
- **Print:** Print the wiki page.

Additional features appear below the wiki page's content. A view counter displays the wiki page's view count. Ratings and comments also appear if they're enabled.

53.4 Wiki Page Details

When viewing a wiki page, you can view its details by clicking *Details* above the page content. Several tabs appear, to give you access to several categories of information about the page.

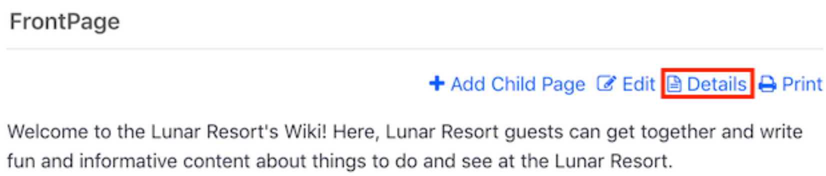

Figure 53.9: Click *Details* to view the wiki page's details.

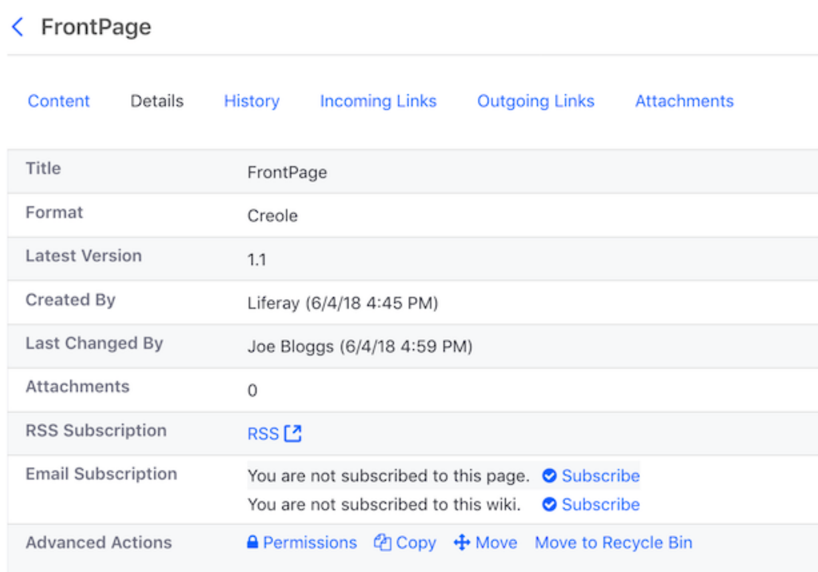

Figure 53.10: The wiki page's details.

Details

The Details tab shows page statistics and lets you perform some actions on the page:

- **Title**: The page title.
- **Format**: The page's format (Creole, HTML, MediaWiki, or plain text).
- **Latest Version**: The page's latest version. The wiki automatically tracks page versions whenever a page is edited.
- **Created By**: The user who created the page.
- **Last Changed By**: The user who last modified the page.
- **Attachments**: The number of attachments to the page.

- **RSS Subscription**: An icon that opens a new page where you can subscribe to an RSS feed using Live Bookmarks, Yahoo, Microsoft Outlook, or an application you can choose from your machine.
- **Email Subscription**: Links that let you to subscribe to or unsubscribe from modifications notifications for the page and the entire wiki node.
- **Advanced Actions**: Links that let you modify the page's permissions, make a copy of the page, move (rename) the page, or move the page to the recycle bin.

History

The History tab lets you access the page's activities and versions via tabs:

- **Activities:** Lists actions performed on the page. Each activity has an icon that represents the type of action, the name of the user, the action's description, date, and an *Actions* menu (⋮) to revert the action or compare its resulting version to that of another action.

- **Versions:** Lists all the wiki page's versions. You can revert a page back to a previous version by selecting *Revert* from that version's *Actions* menu (⋮). You can also compare the differences between versions by selecting two versions and then clicking the *Compare Versions* button.

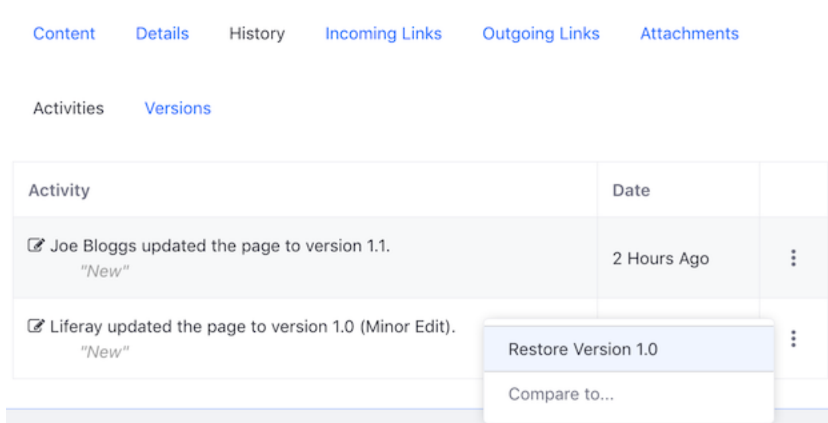

Figure 53.11: The Activities tab displays the actions taken on the wiki page.

Incoming/Outgoing Links

The tabs *Incoming Links* and *Outgoing Links* list incoming and outgoing links, respectively. These are wiki links to and from the wiki page. You can use this tab to examine how this page links to other pages and how other pages link back to this page.

Attachments

The *Attachments* tab lists the name and size of each file attached to the page. You can attach any file to the wiki. Images are the most common type of file attached to a page. Referencing them using the proper WikiCreole syntax renders the image inline, which is a nice way to include illustrations in your wiki documents.

53.5 Other Wiki Widgets

The widgets that accompany the main Wiki widget help you display and navigate particular wiki nodes. The following widgets are available:

- **Page Menu:** Displays a single wiki page's outgoing links.
- **Tree Menu:** Displays a wiki's page hierarchy as a tree.
- **Wiki Display:** Displays a single wiki node.

You can find these widgets in the *Add* () → *Widgets* → *Wiki* menu.

Page Menu

The Page Menu widget displays a wiki page's outgoing links. It answers the question, "What wiki pages can I access from this page?" After adding the Page Menu widget to a site page, you must set the wiki page it displays links from. Follow these steps to do so:

1. Click the widget's *Options* icon ( ) and select *Configuration*.
2. In the configuration's *Setup* tab, choose the wiki node then click *Save*.
3. Still in the configuration's *Setup* tab, select the wiki page then click *Save* and close the configuration dialog.

When you click a Page Menu link, the site page's Wiki or Wiki Display widget displays the wiki page associated with that link.

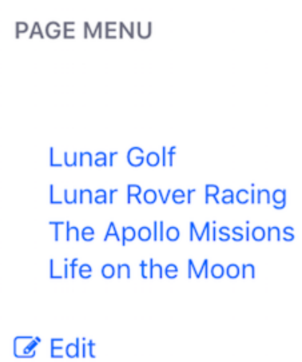

Figure 53.12: The Page Menu widget displays a wiki page's outgoing links.

Tree Menu

The Tree Menu widget displays a wiki's page hierarchy as a tree that lets you navigate all the wiki's pages. Much like the Page Menu setup, you configure the Tree Menu widget to focus on a wiki node. You can also configure how deep users can navigate into the page hierarchy. You can set the *Depth* to a value from 1 to 5, or select *All* to allow navigation to all of the wiki node's pages.

Follow these steps to configure the Tree Menu widget after adding it to a site page:

1. Click the widget's *Options* icon ( ) and select *Configuration*.
2. In the configuration's *Setup* tab, choose the wiki node then select the depth of wiki pages to display in the hierarchy.
3. Click *Save* and close the configuration dialog.

In the Tree Menu, folder icons represent parent wiki pages and document icons represent child wiki pages at the end of the nodes. When you click a parent wiki page or child wiki page, the Wiki or Wiki Display widgets on the site page display the respective wiki page.

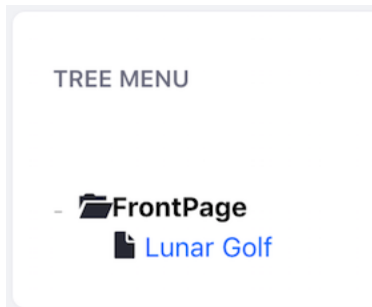

Figure 53.13: The Tree Menu widget displays a wiki node's hierarchy to the configured depth.

Wiki Display

The Wiki Display widget lets you focus user attention on one wiki node. After adding the widget to a page, follow these steps to configure it:

1. Click the widget's *Options* icon ( ) and select *Configuration*.
2. In the configuration's *Setup* tab, choose the wiki node then click *Save*.
3. Still in the configuration's *Setup* tab, select the wiki page then click *Save* and close the configuration dialog. This page serves as the entry point for the wiki.

The configuration options and user interface for the Wiki Display are almost identical to that of the Wiki widget.

53.6 Sending Alerts and Announcements

You can use the Alerts and Announcements widgets on Site pages to broadcast important information to users. The Alerts widget is designed for displaying high-priority information (e.g. planned downtime alerts, security alerts, etc.). Each alert is therefore labeled with a red *Important* tag. The Announcements widget displays all other information you want to broadcast on your site. Each announcement therefore lacks the red tag. To separate important alerts from more mundane announcements, you can place the Alerts and Announcements widgets on different pages. However,

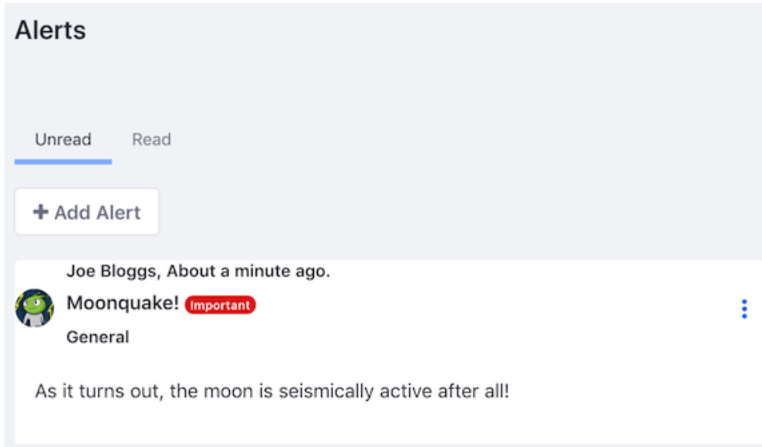

Figure 53.14: The Alerts widget provides administrators with an easy way to communicate important information to appropriate groups of users.

you can use either widget to display any information you wish. Besides the red tag, they function the same. You can also scope your alerts and announcements to specific groups of users.

These widgets have two tabs:

- **Unread:** Non-expired alerts/announcements that you haven't read.
- **Read:** Alerts/announcements that have expired, or that you've read.

Click an alert/announcement's *Actions* button (⋮) to edit or delete it.

Creating Alerts and Announcements

There are two places where you can create alerts and announcements:

1. The *Announcements and Alerts* app in the Control Panel. In this app, announcements and alerts are in separate tabs. To begin creating an announcement or alert, select the appropriate tab and then click the *Add* button (+). This app gives administrators a central location to create announcements and alerts that are then displayed on site pages by the Announcements and Alerts widgets.

2. The Announcements and Alerts widgets, after adding them to a site page from the *Add* (+) → *Widgets* → *News* menu. To begin creating an announcement or alert, click the widget's *Add Alert* or *Add Announcement* button.

Regardless of where you create the alert or announcement, the form for creating it is the same. Follow these steps to complete the form:

1. Use the *Title* field to give the alert or announcement a title. Then create your content in the field *Write your content here....* For a detailed explanation of the editor, see the Blogs documentation.

2. Open the *Configuration* section of the form and set the following options, if desired:

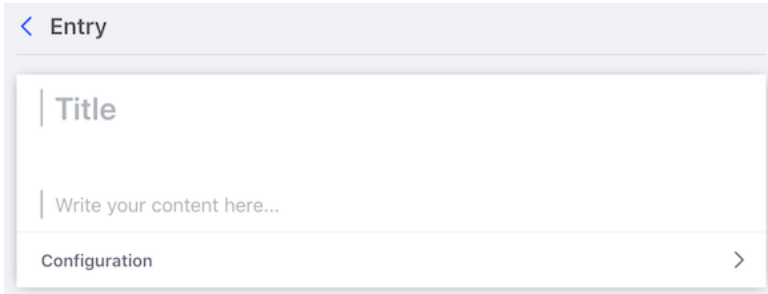

Figure 53.15: Enter your alert or announcement's title and content.

- **Distribution Scope:** The scope where the alert/announcement is displayed. The default *General* scope sends the alert/announcements to everyone. Alternatively, you can select your site or specific roles as the scope.
- **URL:** A URL (optional) to include with the alert/announcement. For example, an announcement about a news story could include a link to the news article. The URL must be valid and begin with `http://` or `https://`.
- **Type:** The alert/announcement type. This can be *General*, *News*, or *Test*. Note that each user can specify a different delivery mechanism for each type of alert/announcement. See the User Configuration section for details.
- **Priority:** The announcement's priority. This can be *Normal* or *Important*. Note that this is disabled for alerts because alerts are always high priority.
- **Display Date:** The display date of the alert/announcement. This determines when the alert/announcement is sent to users and appears in the widget. By default, the *Display Immediately* box is checked. This sets the display date equal to the creation date. Uncheck this box to enter a custom display date. For example, administrators can create alerts/announcements for display on a later date. This date can be days, weeks, months, or years in the future. Once the *Display Immediately* box is unchecked, clicking the Display Date field opens the date-picker.
- **Expiration Date:** The date and time the alert/announcement expires. Once an alert/announcement expires, the widget displays it in the Read tab. Clicking the Expiration Date field opens the date-picker.

3. Click *Save* when you're done. Your alert/announcement then appears in the widget.

User Configuration

Users can configure how they'd like to receive announcements.

1. From the Menu (), open your user menu and select *My Account* → *Account Settings*.
2. On the *Preferences* tab, select *Alerts and Announcements Delivery*. This brings up a menu for customizing the delivery options for alerts and announcements.
3. Select a configuration for each type of alert/announcement (General, News, or Test). For each type, you can enable delivery by email and SMS (text message). Note that the *Website* delivery option is selected and grayed out for each alert type. This means that each alert/announcement is always viewable in its respective widget on a site.

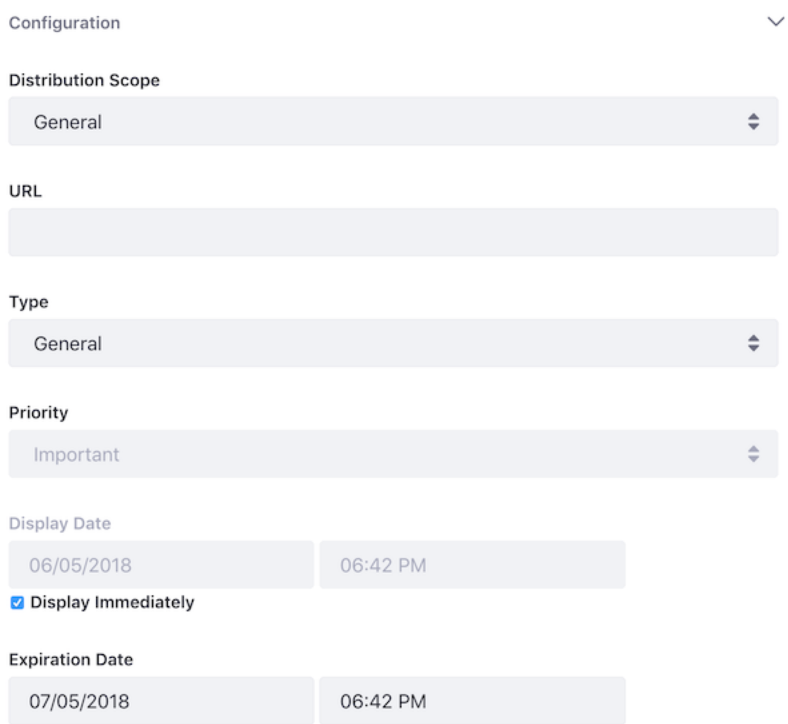

Figure 53.16: Configure your new alert or announcement.

4. Click *Save* when you're finished.

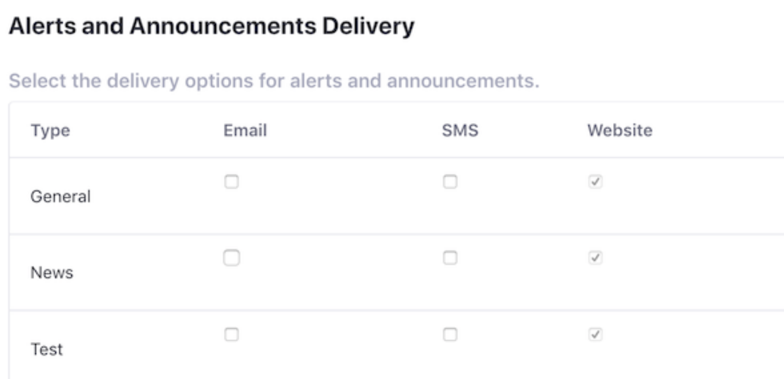

Figure 53.17: Each user can choose how they receive alerts and announcements.

Alert and Announcement Roles

You can also create roles for users to make general announcements. For instance, if you want someone specific to have strict control over announcements, give that person an Announcements Role. Follow these steps to create a simple Announcements Role:

1. Navigate to *Control Panel* → *Users* → *Roles*.

2. With the *Regular Roles* tab selected, click the *Add* button (+). This opens the *New Role* form.

3. Name your Role *Announcements*, give it a title and description, and click *Save*.
4. Select the Role's *Define Permissions* tab then grant these permissions:
    - In *Control Panel* → *General Permissions*, select *Add General Announcements*.
    - In *Site Administration* → *Applications* → *Announcements*, select all the resource permissions.

    Click *Save* after selecting each permission. These permissions let the Role add alerts and announcements.

Now you have a simple Announcements Role that can manage your site's general announcements. Of course, you can adjust this Role's permissions.

53.7 Managing Notifications and Requests

If you subscribed to a blog or message board, or if someone sent you a private message, invitation, event reminder, or mentioned you in a post, you received a notification or request.

The notifications icon on your profile image shows your number of unread notifications or requests.

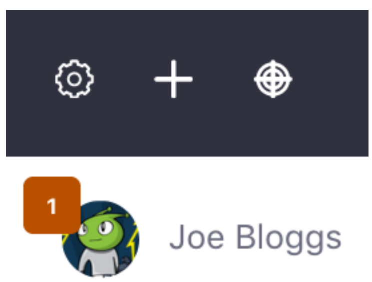

Figure 53.18: The number of notifications and requests are displayed above your profile image

To access notifications and requests, click the notifications icon on your user profile image or open the Menu () and select *My Account* → *Notifications* under your user menu. The *Notifications List* tab is selected by default. This is where all your notifications appear. Click the *Requests List* tab to view and manage your requests.

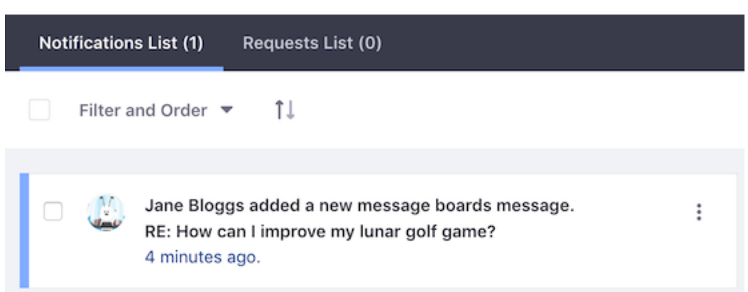

Figure 53.19: The *Notifications List* section displays all your notifications in a paginated list.

Managing Notifications

Notifications can pile up after some time, especially if you were away for a few days. The Management Bar gives you several ways to filter and sort your notifications.

The *Filter and Order* menu gives you the following options for viewing notifications:

- **All:** The default option. Displays both read and unread notifications.
- **Unread:** Displays notifications that haven't been marked as read. Unread notifications are indicated with a blue border on the left-hand side of the notification.
- **Read:** Displays notifications that have been marked as read.
- **Date:** Order notifications by date.

By default, notifications are listed by date in descending order. To sort notifications by ascending order, click the up/down arrow icon in the management bar. Clicking the button again reverses the sort.

Each notification's *Actions* menu (⋮) lets you mark the notification as read/unread, or delete the notification.

Managing Multiple Notifications

You can also manage multiple notifications at once. Select the checkbox next to notifications you want to manage and choose an option from the Management Bar. Select the checkbox above the notifications list to select all notifications on the current page. The Management Bar shows three actions for selected notifications:

- Mark as Read ()
- Mark as Unread ()
- Delete (✕)

You can also select *Mark all notifications as read* in the Options menu (⋮). Note that this feature is only available if at least one notification is unread. Clicking it marks all notifications as read, selected or not.

Managing Requests

When you get a request, it appears in the *Requests List* tab. In each request's *Actions* menu (⋮), you can click *Confirm* to accept, *Ignore* to decline, or *Delete* to remove the request.

CHAPTER 54

USING THE KNOWLEDGE BASE

The Knowledge Base app can be used to display professional product documentation or form complete books or guides. It even lets you import article source files written in Markdown. It's workflow-enabled, so you can require articles to be approved before publication. Additionally, you can create article templates that help users follow a common outline.

Note: To use the Knowledge Base, you must download and install the Knowledge Base app from Liferay Marketplace.

Here are the Knowledge Base's key features:

- Navigation is built into Knowledge Base Display.
- The suggestions interface enables user feedback on articles.
- Articles are stored in folders.
- Metadata fields exist for the friendly URL, source URL, categorization, and related assets.
- The *Edit on GitHub* button (Edit on GitHub) can take readers to an article's source repository location (if you choose to use it that way).
- Markdown source files can be imported to create and update articles.

The Knowledge Base has several widgets you can add to site pages:

- Knowledge Base Article
- Knowledge Base Display
- Knowledge Base Search
- Knowledge Base Section

54.1 Creating Knowledge Base Articles

The Knowledge Base app in Site Administration contains everything you need to create articles in the Knowledge Base. You can create articles by authoring them in the app's WYSIWYG editor or by importing them from Markdown files (`.markdown`, `.md`) in a ZIP archive. The sections below cover both ways of creating articles.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base* → *Access in Site Administration*. To add or act on articles, folders, or suggestions, the Site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

To navigate to the Knowledge Base app, open the Menu () then go to Site Administration (the menu for your Site) → *Content* → *Knowledge Base*. The Knowledge Base app has 3 tabs:

- **Articles:** Create and manage articles and folders.
- **Templates:** Create and manage templates.
- **Suggestions:** Manage user-submitted feedback for articles.

Select the *Articles* tab, then proceed to the sections below for instructions on creating articles.

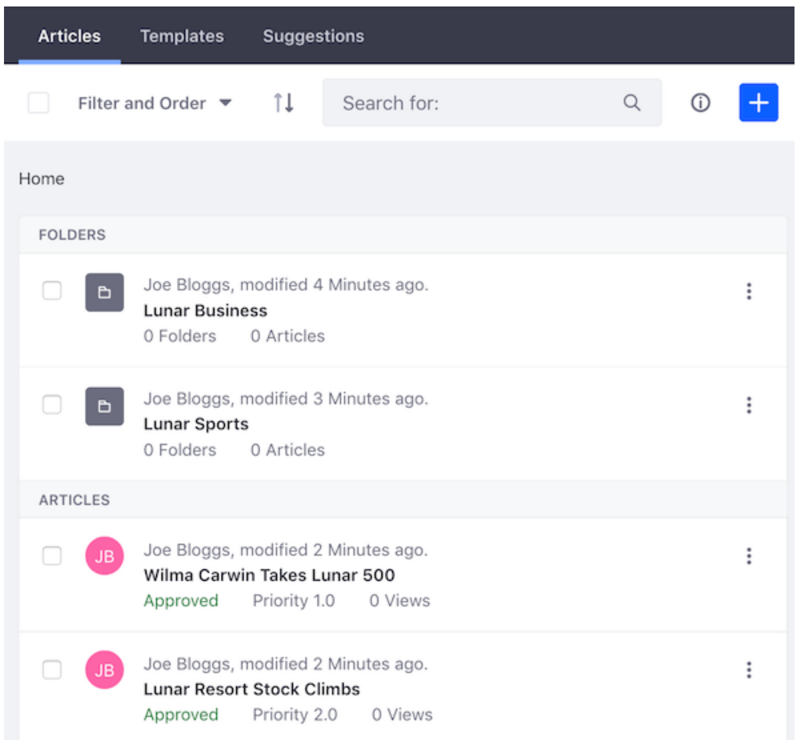

Figure 54.1: The Knowledge Base app in Site Administration lets you create Knowledge Base articles.

Authoring Articles in the Editor

Follow these steps to create an article in the editor:

1. In the Articles tab, click the *Add* button () and choose *Basic Article* or the name of an available template. This brings up the New Article form.

2. Enter a title for the article. A URL-safe version of the title you provide is added to the end of the article's friendly URL. You can manage the friendly URL in the *Configuration* section's *Friendly URL* field.

3. Use the WYSIWYG editor to create the article's content. To view or edit the article's HTML source, click the *Source* button in the editor. The sections below the editor let you add attachments and tags, specify related assets, and set permissions for the article. By default, View permission is granted to the Guest role, meaning anyone can view your article.

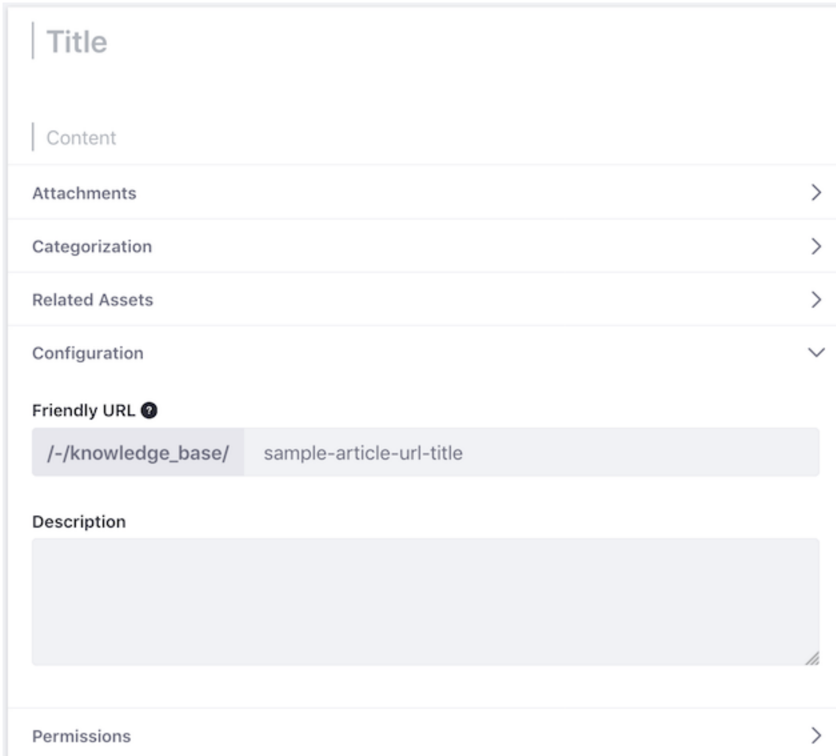

Figure 54.2: You can create and modify a Knowledge Base article's content using the WYSIWYG editor.

4. Click *Publish* to submit the article for publication or click *Save as Draft* to continue working on it later. Note that if you've enabled workflow for the Knowledge Base, your article must be approved before publication.

Once the article is saved, it is converted automatically to HTML for the Knowledge Base. Articles are listed in a table in the Articles tab.

Importing Knowledge Base Articles

You can also create new Knowledge Base articles by importing them from a ZIP archive that contains articles in the Markdown format (`.markdown`, `.md`). For example, you could write articles in your favorite Markdown editor, package them in a ZIP file, and then import that ZIP file to create those articles in the Knowledge Base. The Knowledge Base can also prioritize articles by their filenames' numerical prefixes. For example, the Knowledge Base would list `01-article.markdown` and `02-article.markdown` in ascending order by their numerical prefix (`01`, `02`). For more information on article priority, see Managing Knowledge Base Articles For detailed information on the Knowledge Base importer, see the following topics:

- Importing Knowledge Base Articles

- Knowledge Base ZIP File Requirements
- Knowledge Base Importer FAQs

Note: To import articles, your Role must have the permission *Knowledge Base → Resource Permissions: Import Articles*.

Follow these steps to import articles into the Knowledge Base:

1. In the Articles tab, click the *Add* button () and choose *Import*. This brings up the Import form.

2. Click *Browse* and select the ZIP file that contains the articles you want to import.

3. If you want to use the files' numerical prefixes to prioritize the imported articles in the Knowledge Base, select *Apply numerical prefixes of article files as priorities*.

4. Click *Save* when you're finished.

Like all articles, imported articles are automatically converted to HTML for the Knowledge Base and listed in a table with the rest of the articles in the Articles tab.

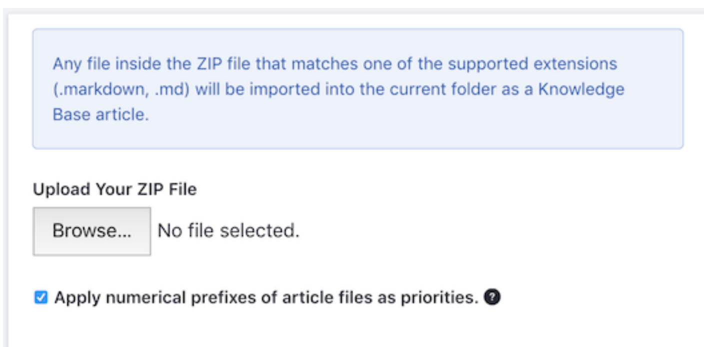

Figure 54.3: You can import ZIP files that contain Knowledge Base articles in Markdown format.

54.2 Managing the Knowledge Base

The Knowledge Base app in Site Administration manages the Knowledge Base. To navigate to this app, open the Menu () then go to Site Administration (the menu for your site) → *Content* → *Knowledge Base*.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base → Access in Site Administration*. To add or act on articles, folders, or suggestions, the site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

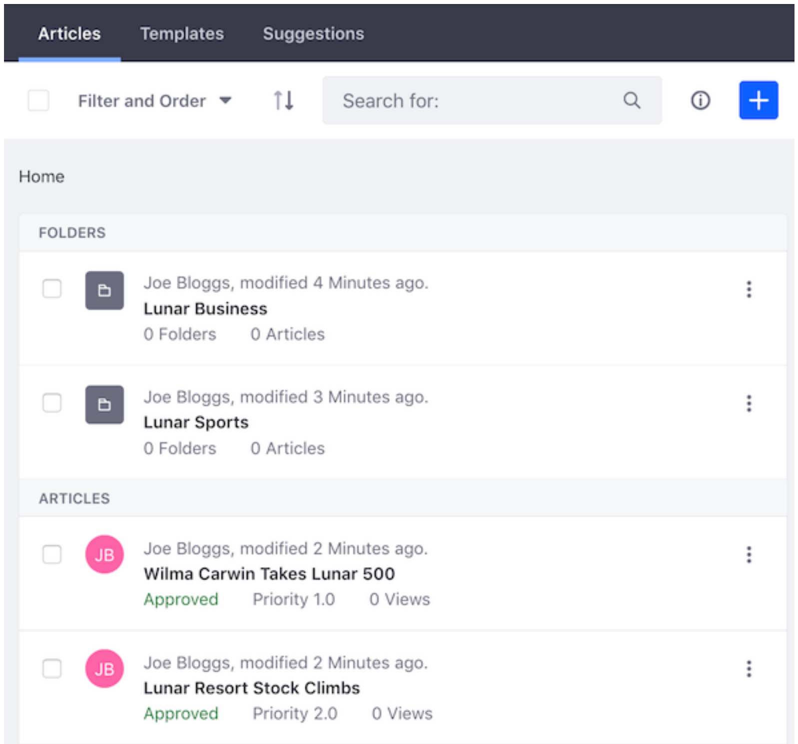

Figure 54.4: You can manage Knowledge Base articles, folders, and suggestions.

Setting the Knowledge Base's Options

At the top-right of the Knowledge Base app, the Options menu () contains these options:

Subscribe: Get notified when Knowledge Base articles are created, updated, or deleted.

Home Folder Permissions: Define detailed permissions for the Knowledge Base app. You can choose the Roles that can perform the following tasks:

- Add/delete articles, folders, and templates
- Change the Knowledge Base app's permissions
- Subscribe to articles
- View templates and suggestions

Export/Import: Export or import the Knowledge Base app's configuration.

Configuration: Configure email notifications for article subscriptions and suggestions. You can also make the Knowledge Base app's articles available via RSS (enabled by default), and configure the RSS feed's options.

Managing Knowledge Base Articles

Each article also has a *priority* value that determines its position in the Knowledge Base Display widget's navigation. Each article's priority value appears beneath the article's title. The Knowledge Base Display widget's navigation arranges articles in ascending priority. Priority 1 is the highest priority. The higher an article's priority, the higher it appears in the navigation. Articles are assigned the next lowest priority by default. This behavior can be changed via Knowledge Base System Settings.

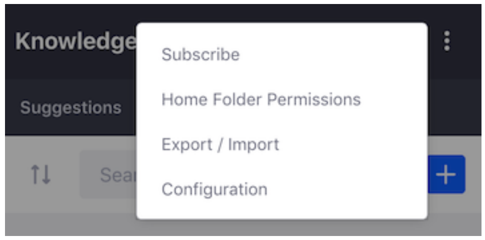

Figure 54.5: The Knowledge Base App's options.

To assign articles a new priority value, follow these steps:

1. Select *Move* from the Actions menu (⋮) next to the article.
2. Enter a new priority value for the article.
3. Click *Move* to apply the new priority.

You can also organize articles with folders. Follow these steps to create a folder:

1. Click the Add button (+) and select *Folder*. This opens a form for creating the new folder.
2. Enter a name and an optional description.

   By default, anyone can view the folder. You can manage this setting along with the other permissions in the form's *Permissions* section.
3. Click *Save*. The folder is then listed in a table in the Articles tab.

The text immediately below the *Filter and Order* selector at the top of the app shows your position in the folder hierarchy. Click a folder's name in the hierarchy to navigate to it. You can also move articles into folders and create child articles. Knowledge Base also supports nested folders.

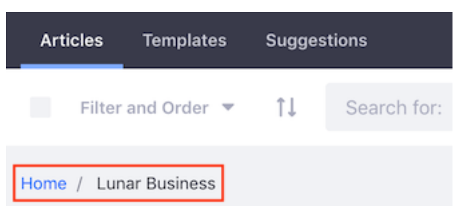

Figure 54.6: This screenshot uses a red box to highlight the text that indicates the current position in the folder hierarchy.

Each folder's Actions menu (⋮) lets you perform the following actions on the folder:
Edit: Change the folder's name and description.
Move: Relocate the folder under a new parent folder or update its priority.
Delete: Remove the folder and its articles from the Knowledge Base.
Permissions: Grant or revoke the following permissions: add an article to the folder, add a sub-folder to the folder, delete the folder, move the folder, set permissions on the folder, edit (update) the folder, and view the folder.

You can also delete multiple articles or folders at once. To do this, select the checkbox for each item that you want to delete and click the *X* button that appears in the Management Bar. You can also see the info for selected items by clicking the *Info* button (ⓘ) in the Management Bar.

54.3 Knowledge Base Templates

Templates give users a starting point. For example, you can create templates that contain default headers or other content for articles. Templates help foster consistent formatting and content organization for articles. You can create and manage templates from the Knowledge Base app in Site Administration. To navigate to this app, open the Menu () and go to Site Administration (the menu for your Site) → *Content* → *Knowledge Base*.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base* → *Access in Site Administration*. To add or act on articles, folders, or suggestions, the Site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

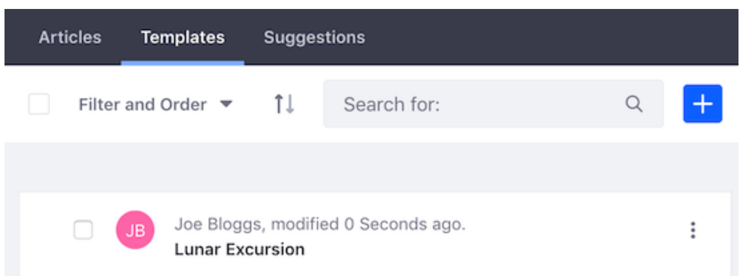

Figure 54.7: The Knowledge Base app's Templates tab.

Creating Templates

To create a new template, follow these steps:

1. Click the *Templates* tab.
2. Click the *Add* button (). This brings up the *New Template* form.
3. Enter a title for the template.
4. Use the WYSIWYG editor to create the template's content. To view or edit the article's HTML source, click the *Source* button (</>) in the editor. You can also set the template's permissions via the form's *Permissions* section.
5. Click *Publish* to finish creating the template.

Figure 54.8: The New Template form.

Managing Templates

Each template appears in a list in the Templates tab. You can take the following actions on each template via its Actions button ( ):

View: Display the template. From here, you can print the template, use it to create an article, edit it, modify its permissions, or delete it.

Edit: Change the template's title and content.

Permissions: Configure the template's permissions. You can choose whether a Role can change permissions, update, view, or delete the template.

Delete: Remove the template from the Knowledge Base.

54.4 Responding to Knowledge Base Feedback

The Knowledge Base app's *Suggestions* tab shows user feedback on articles and lets you mark progress on addressing the feedback. To navigate to this app, open the Menu ( ) then go to Site Administration (the menu for your Site) → *Content* → *Knowledge Base*.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base* → *Access in Site Administration*. To add or act on articles, folders, or suggestions, the Site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

Each suggestion provides the link to the associated article, the user's feedback, the user's name, the feedback's time stamp, and the status on addressing the suggestion. You can use each entry's Actions menu ( ) to move the entry between the *New*, *In Progress*, and *Resolved* states.

Note: To view article suggestions, your Role must have the permission *Knowledge Base* → *Knowledge Base: View Suggestions*.

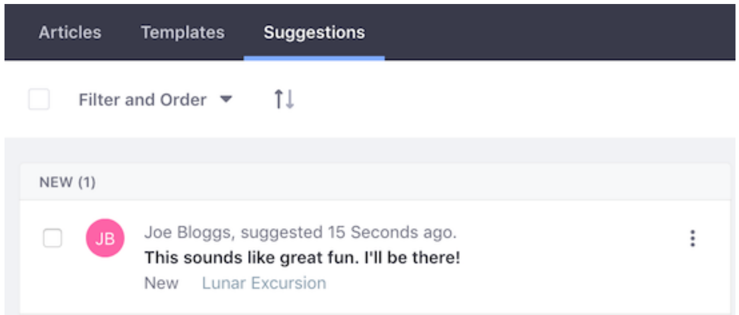

Figure 54.9: The Suggestions tab in Knowledge Base displays each piece of feedback that users leave on Knowledge Base articles.

To move suggestions between the *New*, *In Progress*, and *Resolved* states, your Role must have the permission *Knowledge Base → Knowledge Base Article: Update*. Roles assigned this permission can also view and update the state of article suggestions from any of the other Knowledge Base widgets.

When you move the suggestion to a different state, an email is sent notifying the user of the change. You can view and configure the automated emails from the Knowledge Base app's *Options* ( ) → *Configuration* menu.

54.5 Knowledge Base Display

You can use the Knowledge Base Display widget to display your published Knowledge Base articles. You can customize how this widget displays articles, and which ones it displays. To get started, add the widget to the Site page you want to display articles on:

1. Navigate to the page and open the *Add* menu ( ).
2. Open the *Widgets → Content Management* section, then add *Knowledge Base Display* to the page.

By default, the Knowledge Base Display widget displays articles from the Knowledge Base's Home folder. To change the location of its articles, follow these steps:

1. Click the widget's *Options* icon ( ) and select *Configuration*. This opens the Configuration dialog.
2. In the *Setup* tab, select the *General* tab then click *Select* in the *Article or Folder* field. This brings up the *Select Entry* form.
3. Click *Choose* next to the article or folder of articles you want to display. Alternatively, you can click the *Choose This Folder* button at the top of the Select Entry form to select the current folder.
4. Click *Save* and close the Configuration dialog.

The Knowledge Base Display widget's Options icon ( ) also provides these common configuration options:

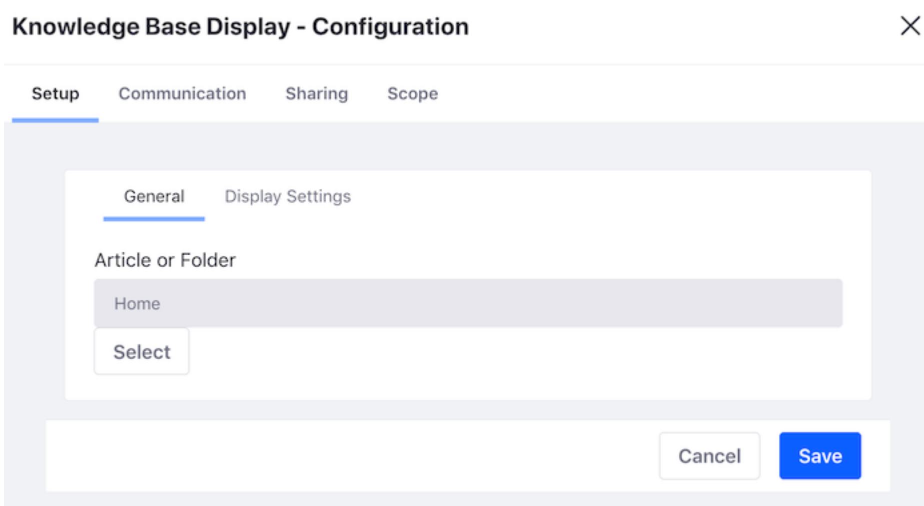

Figure 54.10: Select the article or folder of articles that the Knowledge Base Display widget displays.

- Look and Feel Configuration
- Export/Import
- Permissions
- Configuration Templates

For more information on these, see the section on configuring widgets in Web Experience Management.

The Knowledge Base Display's navigation menu and display options make it the perfect candidate for a full page widget. If you display a folder of articles, the navigation on the left side of the widget displays links to all the folder's articles. The viewing area on the right side of the widget displays the folder's leading article (the *priority one* article). Click an article in the navigation to display it in the viewing area. The currently displayed article's link appears in bold in the navigation. You can also move between articles by clicking the links with arrows at the bottom of the widget.

Knowledge Base Display can also show article hierarchies. Viewing an article that has child articles expands the navigation tree to show links to the child articles. Any expanded nodes collapse when you view a different top level article.

The links at the top of the widget allow users to perform the following actions on an article:

- Subscribe to an RSS feed of the Knowledge Base
- Subscribe to the current article
- View the current article's history
- Print the current article

Administrators have access to an additional set of links at the top of the widget that lets them perform the following actions:

- Edit the article
- Add a child article
- Set the article's permissions
- Move the article
- Delete the article

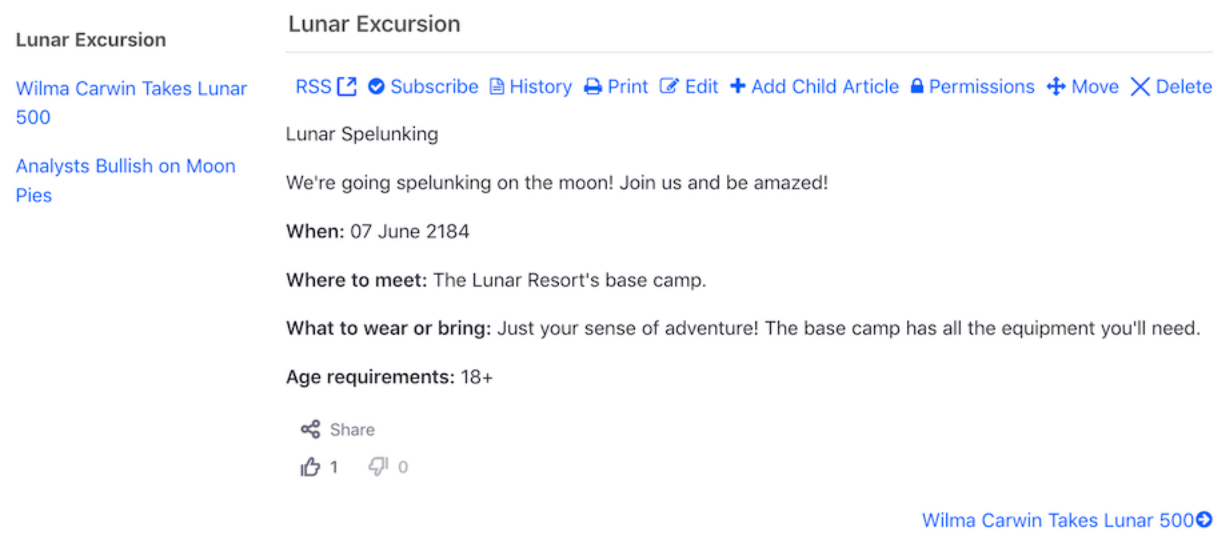

Figure 54.11: Knowledge Base Display's navigation and viewing provide a great reading experience.

Below the article's content is the rating interface, showing thumbs up/down icons. Users can also submit suggestions or comments below the article in the text box labeled *Do you have any suggestions?*. Administrators can view the suggestions and mark progress on them.

If the administrator enables the Knowledge Base app's source URL feature (more on this in a moment) and an article has an assigned source URL, an *Edit on GitHub* button (Edit on GitHub) appears to the right of the article's title. This button lets users access the article's source in GitHub. You can use this feature to encourage users to contribute fixes or improvements to articles. If you're interested in this feature, you can direct your administrator to follow the instructions in Knowledge Base System Settings.

Displaying Different Article Sets

As an administrator, say that you've used folders to aggregate similar articles, and you want to provide an easy way for users to switch between these sets of articles. The Knowledge Base Display's content folder feature adds a selector to the top of the navigation that lets users switch between article sets.

Follow these steps to set up content folders:

1. Add a folder in the Knowledge Base app in Site Administration. Then create sub-folders in this folder. These sub-folders are the *content folders*.

2. Add articles to each content folder.

3. Select *Configuration* from Knowledge Base Display's *Options* menu ( ). In the *Setup* → *General* tab, select the content folders' parent folder and click *Save*.

The content selector's values reflect the names of your content folders. Select one to view its articles.

You can also add a common prefix to the names shown in the selector:

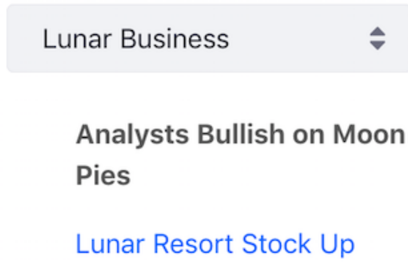

Figure 54.12: Knowledge Base Display's content folder feature lets users switch between different sets of articles.

1. Select *Configuration* from Knowledge Base Display's *Options* menu ( ). In the Configuration dialog, select the *Setup* → *Display Settings* tab.

2. Enter the prefix into the *Content Root Prefix* field and click *Save*.

54.6 Other Knowledge Base Widgets

There are other Knowledge Base widgets you can add to Site pages besides Knowledge Base Display:

- **Knowledge Base Article:** Display a single article's content.
- **Knowledge Base Section:** Publish articles associated with a specific topic (section).
- **Knowledge Base Search:** Search for articles.

You can add these widgets from *Add* ( ) → *Widgets* → *Content Management*.

Knowledge Base Article

Knowledge Base Article displays a single article's content. It even shows abstracts of child articles. You can add multiple Knowledge Base Article instances to a page, and each one can show a different article.

After adding Knowledge Base Article to a page, follow these steps to configure the widget:

1. Click *Please configure this portlet to make it visible to all users*. This opens the Configuration dialog.

2. In the *Setup* → *General* tab, click *Select*, choose an article, click *Save*, and close the Configuration dialog.

You can change your selection from the widget's *Options* ( ) → *Configuration* menu.

Knowledge Base Article shares the same UI as the Knowledge Base Display to display and manage its articles. Refer to the Knowledge Base Display documentation for a detailed description of the widget's UI.

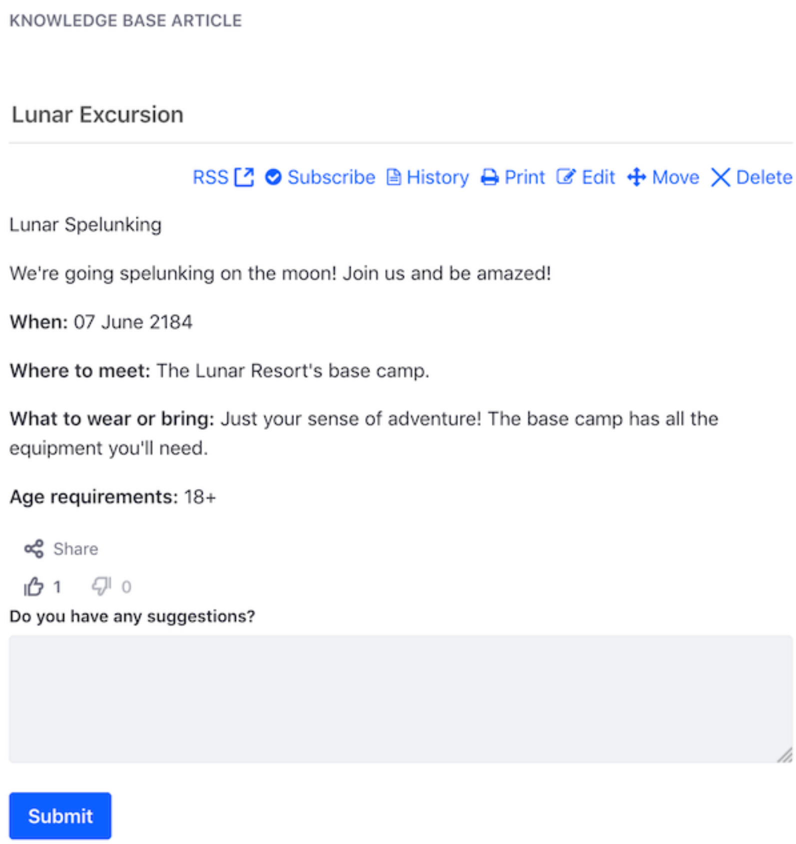

Figure 54.13: The Knowledge Base Article app is great at displaying individual articles.

Knowledge Base Section

Note: as of Knowledge Base 3.0.0, the Knowledge Base Sections widget is deprecated and replaced by categories.

The Knowledge Base Section widget lets you publish articles associated with a specific topic (section). For example, a news Site might have the sections *World*, *Politics*, *Business*, and *Entertainment*.

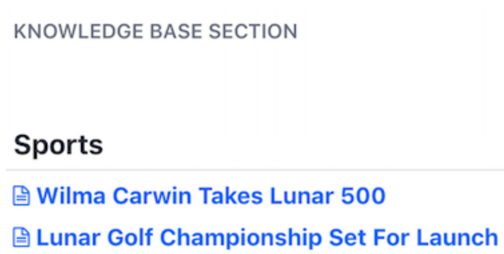

Figure 54.14: The Knowledge Base Section widget.

To use sections, an administrator must first configure the feature in System Settings, creating the section names for use in the Knowledge Base Section widget. This process is covered in detail in Knowledge Base System Settings. When creating or editing a Knowledge Base article, authors can then select the article's section in the *Configuration* → *Section* field.

You can add multiple instances of the Knowledge Base Section widget to a page. Each widget can display articles from any number of sections. You can configure the widget to display article titles or abstracts. You can also define whether to show pagination or section titles.

Follow these steps to configure an instance of the Knowledge Base Section widget:

1. Select *Configuration* from the Knowledge Base Section widget's *Options* menu (⋮). This opens the widget's Configuration dialog.

2. In the Configuration dialog's *Setup* → *General* tab, select the section or sections that you want to use and click *Save*.

3. Close the Configuration dialog to see the updates.

The matching articles are displayed in the app beneath their section heading.

Knowledge Base Search

Note: as of Knowledge Base 3.0.0, the Knowledge Base Search widget is deprecated and replaced by Liferay Search.

Even though the Knowledge Base can show the structure of its articles, it may be difficult to find exactly what you're looking for by browsing. That's where the Knowledge Base Search widget comes in.

Enter your search term and press the *Search* button. The results are displayed in a table with the following criteria for each matching article:

- Title
- Author
- Create date
- Modified date
- Number of views

You can select the criteria to display in the widget's *Options* (⋮) *Configuration* dialog.

54.7 Importing Knowledge Base Articles

As mentioned earlier, the Knowledge Base app can import articles in bulk. This lets you have an offline process where articles are prepared ahead of time before they're published. Articles are imported into the Knowledge Base as Markdown files. Markdown is a text-only file format that is easy to read, yet supports all the things you'd need to do to format your articles.

Note: To import articles, your Role must be granted the *Knowledge Base* → *Resource Permissions: Import Articles* permission.

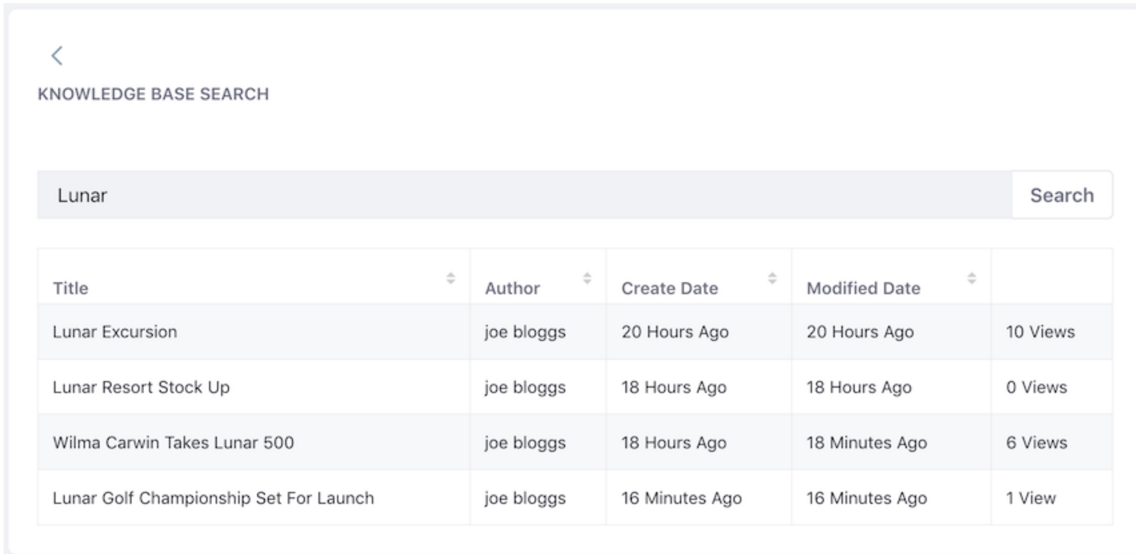

Figure 54.15: The Knowledge Base Search widget lets you search the Knowledge Base for keywords.

The Knowledge Base supports a Markdown dialect known as Multi-Markdown. This dialect extends the original Markdown with features like table formatting, image captions, and footnotes.

For the Knowledge Base to import your Markdown articles, they must adhere to these requirements:

- All source files must use the `.markdown` or `.md` extensions.
- Articles must start with a top-level header (e.g., `# Some Heading ...`).
- Each header must have an associated, unique ID for the article's friendly URL title and for anchor tags in the article's sub headers. Here's an example of a top-level header that correctly specifies an ID:

```
# Some Heading  [](id=some-heading)
```

Here's Markdown source text for a simple example article:

```
# The Moons of Mars [](id=the-moons-of-mars)

As you look up from your chaise lounge, you're sure to see our neighboring
planet Mars. Did you know that Mars has two moons? You might have to break
out a pair of binoculars to see them.

Its two moons are aptly named after the two sons of mythical Roman god Mars.
Their names are Phobos and Deimos.
```

In the first line above, notice the header's ID assignment `id=the-moons-of-mars`. On import, the ID value becomes the Knowledge Base article's URL title.

Markdown is something of a standard: there's Github Flavored Markdown, a proposed common Markdown syntax, forums that support Markdown (reddit, StackExchange, and others), Markdown editors, and an IETF draft for making it an official Internet media type (text/markdown). Why is there so much interest in Markdown?

1. It's readable. Even if you don't know Markdown, you can read it without having to filter out the syntax.

2. It gets out of a writer's way. You don't have to worry about mousing to various icons to change text into a heading or create bulleted lists. Just start typing. The syntax is very intuitive.

3. There are tools to convert it to many other formats, though it was designed to convert to HTML. If your articles are in Markdown, it's straightforward to publish them to the web, mobile formats (Kindle, ePub), and print.

4. Since it's only text, you can use existing tools to collaborate on that text. Using services like GitHub, people can contribute to your articles, and you can see all the changes that have been made to them.

54.8 Knowledge Base ZIP File Requirements

The Knowledge Base importer supports article hierarchies, so Markdown files can be specified anywhere in the ZIP file's directory structure. They can be nested in any number of folders. Image files are the only files supported for attachments.

Note: Imported articles are independent of the workflow settings. This means that **imported articles are automatically approved.**

Only users with the *Import Articles* permission assigned to their Role are able to import articles. This permission can be assigned manually through *Control Panel → Users → Roles*. If you've upgraded from Liferay Portal 6.2, you can also assign this Role to every Role that was already able to add articles with a command from the Gogo shell.

Open the Gogo shell. Type `knowledgeBase:addImportArticlePermissions` and hit enter.

The ZIP file's articles are imported in file order (alphanumerically). To designate an article's priority, add a numeric prefix to its file name. For example, the priorities for articles named `01-file.markdown` and `02-file.markdown` become `1.0` and `2.0`.

To designate an article to be the parent of all other articles in the same source folder, end its file name with `-intro.markdown`. This creates a parent-child hierarchy. You can use the prefix `00` for parent articles to place them at the top of the folder's file order. The importer uses the numeric prefix of an intro file's folder as its article priority.

Here's the underlying logic for the `00` prefix:

- A file prefix of `00` for a non-intro file assigns the resulting article's priority to `1.0`.
- A file prefix of `00` for a top-level intro file sets the article's priority to the first folder numeric prefix found that is `1.0` or greater.

This convention lets you specify priorities for top-level (non-child) articles in your hierarchy.

When importing, keep the checkbox labeled *Apply numerical prefixes of article files as priorities* selected. If a file doesn't have a prefix, its article gets the next available priority (the highest current priority, plus one).

Below is an example ZIP file structure that demonstrates the features mentioned so far:

ZIP File Structure Example:

- `01-winter-events/`
    - `00-winter-excursions-intro.markdown`

    - `01-star-dust-snow-shoeing.markdown`
    - `02-lunar-alpine.markdown`

- `02-summer-events/`
    - `00-summer-excursions-intro.markdown`
    - `01-lunar-rock-scrambling.markdown`
    - `02-extra-terrestrial-mountain-biking.markdown`
    - `03-lunar-olympics/`
        * `00-lunar-olympics-intro.markdown`
        * `01-zero-gravity-diving.markdown`

- `images/`
    - `some-image.png`
    - `another-image.jpeg`

The above ZIP file specifies `00-winter-excursions-intro.markdown` as the parent of its neighboring Markdown files: `01-star-dust-snow-shoeing.markdown` and `02-lunar-alpine.markdown`. Likewise, `00-lunar-olympics-intro.markdown` is the parent of `01-zero-gravity-diving.markdown`. `00-lunar-olympics-intro.markdown` is also the peer of `01-lunar-rock-scrambling.markdown` and `02-extra-terrestrial-mountain-biking.markdown`, and the child of `00-summer-excursions-intro.markdown`.

ZIP Example's Resulting Relationships and Priorities

- `01-winter-events/00-winter-excursions-intro.markdown`
    - **Article:** Winter Excursions
    - **Relationship:** Peer of *Summer Excursions*
    - **Priority: 1.0**
- `01-winter-events/01-star-dust-snow-shoeing.markdown`
    - **Article:** Star Dust Snow Shoeing
    - **Relationship:** Child of *Winter Excursions*
    - **Priority:** 1.0
- `01-winter-events/02-lunar-alpine.markdown`
    - **Article:** Lunar Alpine
    - **Relationship:** Child of *Winter Excursions*
    - **Priority:** 2.0
- `02-summer-events/00-summer-excursions-intro.markdown`
    - **Article:** Summer Excursions
    - **Relationship:** Peer of *Winter Excursions*
    - **Priority: 2.0**
- `02-summer-events/01-lunar-rock-scrambling.markdown`
    - **Article:** Lunar Rock Scrambling
    - **Relationship:** Child of *Summer Excursions*
    - **Priority:** 1.0
- `02-summer-events/02-extra-terrestrial-mountain-biking.markdown`

  - **Article:** Extra Terrestrial Mountain Biking
  - **Relationship:** Child of *Summer Excursions*
  - **Priority:** 2.0

- `02-summer-events/03-summer-olympics/00-lunar-olympics-intro.markdown`
  - **Article:** Lunar Olympics
  - **Relationship:** Child of *Summer Excursions*
  - **Priority:** 3.0
- `02-summer-events/03-summer-olympics/01-zero-gravity-diving.markdown`
  - **Article:** Zero Gravity Diving
  - **Relationship:** Grandchild of *Summer Excursions*
  - **Relationship:** Child of *Opening Ceremonies*
  - **Priority:** 1.0

ZIP files must meet the following requirements:

- Each ZIP file must end in the suffix `.zip`.
- Each ZIP file must contain at least one Markdown source file, optionally organized in folders.
- All referenced image files must be in a folder named `images` in the ZIP file's root.
- Image files must be in a supported format and must use the appropriate file extensions. Supported extensions are `.bmp`,`.gif`,`.jpeg`,`.jpg`, and `.png`. They're specified via an app system setting. For details, see Knowledge Base System Settings.

Once you have your article ZIP file, it's time to import it.
Follow these steps to import your ZIP file:

1. In the Menu ( ), navigate to *Site Administration* (the menu for your site) → *Content* → *Knowledge Base* → *Articles*.
2. Click *Add* ( ) → *Import* to bring up the importer page.
3. Browse to the location of your file, and in most cases leave the checkbox for the article priorities checked, and then click *Save*.

Your file is uploaded, and the importer converts each source file's Markdown text to HTML, applying the HTML to the resulting article. Any image files that are referenced in an article and included in the ZIP file are imported as attachments to the article.

In addition to source files and images, you can configure a base source URL system setting for the importer that specifies your source file's repository location. Each article's *Edit on GitHub* button (if enabled) takes the user to the source location. The importer prefixes each file's path with the base source URL. This constructs a URL to the article's repository source location; it looks like `[base URL]/[article file path]`. Here's an example base source URL:

```
https://github.com/liferay/liferay-docs/blob/master/develop/tutorials
```

The source URL constructed from this base URL and article source file `folder-1/some-article.markdown` would be:

```
https://github.com/liferay/liferay-docs/blob/master/develop/tutorials/folder-1/some-article.markdown
```

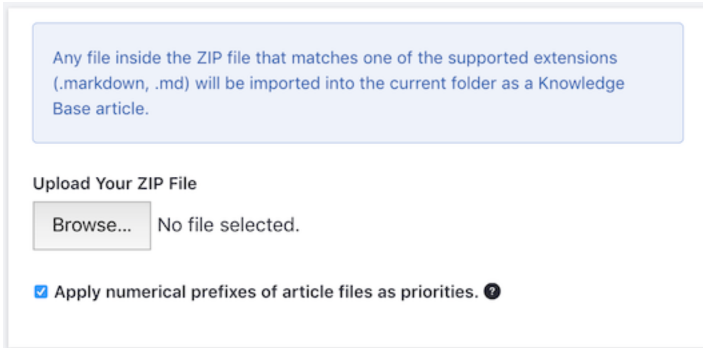

Figure 54.16: Selecting *Add → Import* in Knowledge Base brings up the interface for selecting a ZIP file of Markdown source files and images to produce and update articles in your Knowledge Base.

You specify the base source URL in a file called `.METADATA` in the ZIP file's root folder. The importer treats the `.METADATA` file as a standard Java properties file and uses the base source URL to construct the source URL for all of the ZIP file's resulting articles.

To use the source URL feature, your administrator must enable it via the Knowledge Base System Settings.

54.9 Knowledge Base Importer FAQs

- **What happens when I import an existing article?**

  The importer checks if the source file's leading header ID (e.g., `# Some Heading [](id=some-heading)`) matches the URL title of any existing article in the Knowledge Base folder. If a match is found, the importer replaces the article's content with the incoming content converted from the source file. If no match is found, a new article is created.

- **Do I need to import all of a Knowledge Base folder's articles, even if I only want to create a new article or update a subset of the folder's current articles?**

  No. You can import as many or as few new or modified articles as you like.

- **Does the importer remove articles?**

  No. The importer only creates and updates articles. It doesn't delete any existing articles. To delete an article, you must manually do so via the Knowledge Base app.

- **Can I update an article's priority?**

  Yes. You can use the file/folder prefix convention and re-import the article to update its priority.

- **If I change an article's title, should I also change its header ID?**

  It depends on whether you've already published your article. If it hasn't been published, then there are no public links to it, so it's fine to change the header ID. If the article is already published, you must decide whether it's worth breaking existing links to the article, and worth having search engines rediscover and re-rank your article based on its new friendly URL. The new friendly URL is based on the new header ID.

54.10 Knowledge Base System Settings

Administrators can use the System Settings UI to set the Knowledge Base's global configuration (across sites). You can access this UI in *Control Panel → Configuration → System Settings → Knowledge Base*. There are five sections of Knowledge Base configuration settings:

- Knowledge Base Service
- Knowledge Base Article
- Knowledge Base Display
- Knowledge Base Search
- Knowledge Base Section

The *Knowledge Base Service* section's settings apply defaults to all the Knowledge Base widgets, and to the Knowledge Base app in Site Administration. The other sections apply to specific Knowledge Base widgets and override the *Knowledge Base Service* defaults.

Important: Advanced configuration of the Knowledge Base application's system settings should only be performed by an Liferay DXP administrator.

The Knowledge Base has several optional features that are disabled by default, but can be enabled and configured from System Settings. These include source URL, import file conventions, new article priority increment, and sections.

Source URL Settings

The source URL settings define the source location of Markdown files for import. This should point to a source repository where the files are stored. GitHub is assumed as the default. Once defined, the Knowledge Base displays a button (default label is *Edit on GitHub*) above each displayed article. Users can click the button to navigate to an article's source location.

The source URL settings are accessible in the *Knowledge Base Service* section of the Knowledge Base's System Settings. To enable the source URL, check the *Source URL Enabled* checkbox.

To change the source URL button's label, specify a new value for the setting *Source URL Edit Message Key*. Best practice is to specify the value as a language key. For example, if you create a language key `edit-on-bitbucket=Edit on Bitbucket`, you can specify that language key as the button's new label:

```
edit-on-bitbucket
```

Alternatively, you can specify the label explicitly:

```
Edit on Bitbucket
```

Importer File Convention Settings

These settings define the supported file extensions, the suffix for parent files, and the image folder's path within the imported ZIP files. These settings are accessible in the *Knowledge Base Service* section of the Knowledge Base's System Settings.

The following settings specify the importer's supported file extensions:

Markdown Importer article extensions: Sets the supported article extensions. The default values are `.markdown` and `.md`.

Markdown Importer Image File Extensions: Sets the supported image file extensions. The default values are `.bmp`, `.gif`, `.jpeg`, `.jpg`, and `.png`.

Follow these steps to modify the supported file extensions:

1. Click the + or - button next to the setting to add or remove a supported file extension, respectively.
2. If adding an extension, enter a new value.
3. Click *Save*.

These settings define additional article configuration options for the importer:

Markdown Importer Article Intro: Sets the parent article's file suffix. The default value is `intro.markdown`.

Markdown Importer Image Folder: Sets the image folder path the importer looks for in the ZIP file. The default path is `/images`.

Article Increment Priority Enabled: Whether to increment new article priorities by `1.0`. To disable this increment so that articles get a flat value of `1.0`, deselect the checkbox.

Alternatively, you can enable or disable the article increment priority feature for each widget in the corresponding widget's configuration menu in System Settings.

Section Names Setting

The section names setting lets you specify names of arbitrary topics to attribute to articles. Using the Knowledge Base Section widget, you can display one or more *sections* (groups) of articles. To use sections, you must first define them in the System Settings for the *Knowledge Base Section* widget.

Follow these steps to make new sections available:

1. Navigate to the *Knowledge Base Section* configuration menu.
2. Click the plus button next to the *Admin Knowledge Base Article Sections* setting to add a new field for each section you want.
3. Enter a name for each new section and click *Save*.

Figure 54.17: Create the sections you want to use with the Knowledge Base Section widget.

Once your sections are added, you can follow the steps in the Knowledge Base Section documentation to learn how to use them.

54.11 Bookmarking Sites

The Bookmarks widget provides a list of bookmarks that benefit your Site. For example, a Site that focuses on auto repair could use this widget to provide bookmarks to technical manuals, part specifications, maintenance tips, and more. Administrators can create bookmarks in the widget, and users can access them. Both internal pages as well as external sites can be bookmarked.

Creating Bookmarks and Folders

To get started with the Bookmarks widget, first install it from Liferay Marketplace. It's available for Liferay CE Portal and Liferay DXP via these links:

- Liferay Bookmarks - CE
- Liferay Bookmarks - DXP

See the Marketplace article if you need help using Marketplace to install apps.

After installing the Bookmarks widget, add it to a page. To do this, select *Add* () → *Widgets* → *Community*, and then drag *Bookmarks* onto the page.

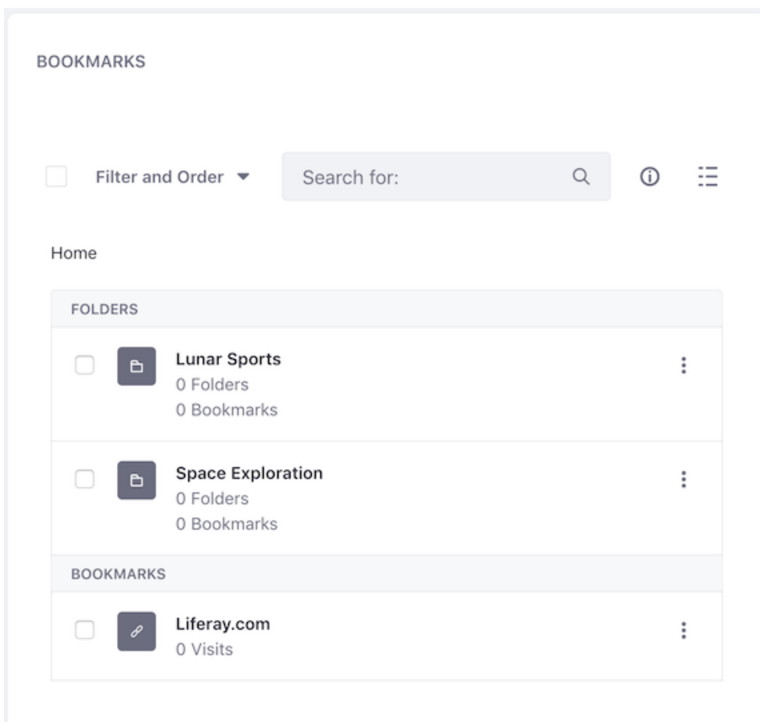

Figure 54.18: The Bookmarks widget displays bookmarks and folders for bookmarks.

The Bookmarks widget is empty at first. The current folder's name (*Home* by default) is displayed to the top-left of the graphic indicating there are no bookmarks in the current folder. The widget displays the Management Bar immediately above the current folder's name. The Management Bar can manage displayed bookmarks. The Management Bar's search bar lets you search for folders and bookmarks in the widget. When you create folders and bookmarks, they're added to the current folder. Folders can have any number of subfolders.

Follow these steps to add a bookmark:

1. In the Bookmarks widget, select *Add* ( ) → *Bookmark*. This opens the Add Bookmark form.

2. Name the bookmark, then enter its URL and description.

3. If you want to associate a tag with the bookmark, open the *Categorization* section and use the *Select* button to select an existing tag. You can also create a new tag by entering it in the *Tags* field and clicking *Add*. See the documentation on tags for more information.

4. If you want to select an existing portal asset (e.g., a media file, blog post, etc.) to relate to the bookmark, open the *Related Assets* section and use the *Select* button to select that asset.

5. Open the *Permissions* section and set the bookmark's permissions.

6. When you're finished filling in the Add Bookmark form, click *Save* to create the bookmark.

You can also create folders for your bookmarks. Follow these steps:

1. In the Bookmarks widget, select *Add* ( ) → *Folder*. This opens the Add Folder form. Note that this form is similar to the Add Bookmark form.

2. Give the folder a name and description.

3. Open the *Permissions* section and set the bookmark's permissions.

4. When you're finished filling in the Add Folder form, click *Save* to create the folder.

Importing and Exporting Bookmarks

Follow these steps to import or export bookmarks:

1. In the Bookmarks widget, select *Options* ( ) → *Export/Import*.

2. In the window that appears, start a new export or import process. For details, see the article on importing and exporting pages and content.

Managing Bookmarks and Folders

The Bookmarks widget provides several ways to manage your bookmarks and folders. One is the Management Bar. Selecting the checkbox on the left-hand side of the Management Bar selects all the folders and bookmarks currently displayed in the widget. When you do this, the Management Bar changes to show the number of selected items. It replaces the buttons on its right with an Info button ( ) and a Recycle Bin button ( ). Pressing the Recycle Bin button moves the selected items to the Recycle Bin. De-select the checkbox to return the Management Bar to its initial state. In this state, the selector menu to the right of the checkbox lets you filter the items to display:

- **All:** Display all folders and bookmarks.
- **Recent:** Display only recent bookmarks.

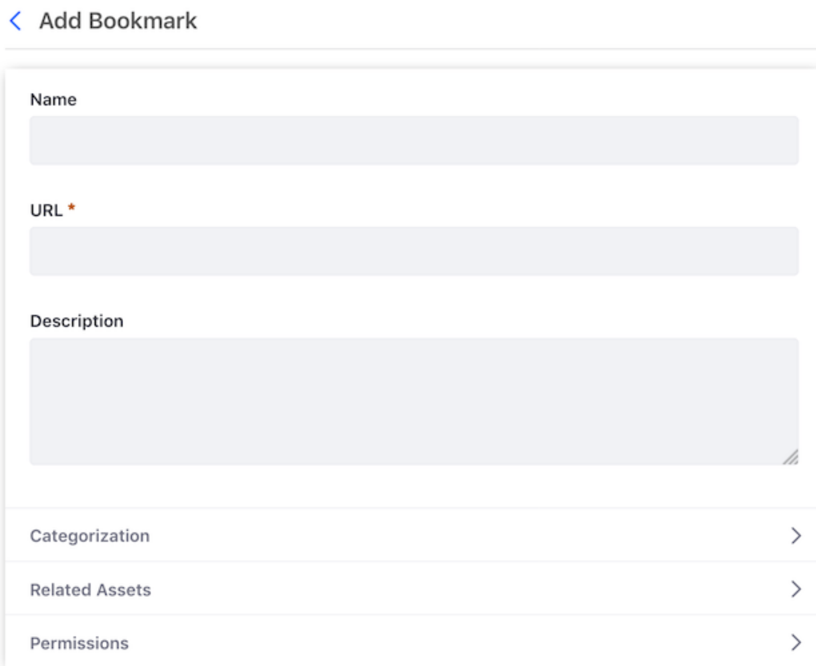

Figure 54.19: The Add Bookmark form lets you create a bookmark.

- **Mine:** Display only the bookmarks you created.

The *View Types* button on the right side of the Management Bar chooses the display of the widget's bookmarks and folders. The View Types button's icon depends on the selected view type:

List ( ): This is the default View Type. It shows items in a simple list that displays the number of bookmarks and subfolders for each folder, and the number of visits for each bookmark.

Table ( ): Shows bookmarks and folders in a table with columns for the name, URL, visits, and modified date of each item.

Regardless of View Type, each item has an Actions menu ( ) menu that lets you perform these actions on the item:

Edit: Edit the item.

Move: Move the item. When moving an item, the Bookmarks widget shows the Move Entries form, where the item's new folder is selected for you automatically. Click *Move* to move the item to this folder. To choose a different folder, click *Select*. The window that appears lets you navigate your folder hierarchy, beginning with the Home folder. Click any folder's name to view its subfolders. The text at the top of the window indicates your position in the hierarchy. You can click the folder names in this text to move back up the hierarchy. To select a subfolder, click its *Choose* button. To choose the folder you're currently browsing, click the *Choose This Folder* button. To create a new folder, click the *Add Subfolder* button. Note that when you're in the Home folder, this is instead the *Add Folder* button. The window closes once you choose a folder, returning you to the Move Entries form. Click *Move* to move the item to the selected folder.

Permissions: View or change the item's permissions.

Subscribe: Subscribe to the bookmark's changes. Note that this action is only available for bookmarks.

Move to Recycle Bin: Move the item to the Recycle Bin.

The Management Bar also contains an Info button (ⓘ) that you can click to slide in an info panel that displays more information about the current folder. This panel also gives you more management options. To receive notifications when a user adds a bookmark to the folder, subscribe to the folder by clicking the *Star* icon (☆). Click the *Actions* menu (⋮) to reveal options for editing, moving, assigning permissions to, or deleting the folder. Note that you can't edit, move, or delete the Home folder.

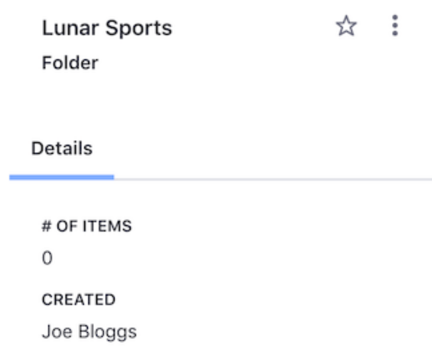

Figure 54.20: A folder's info panel shows more information about the folder, and lets you take additional actions.

54.12 Inviting Members to Your Site

The Invite Members widget lets site administrators send invitations to join the Site. You can add this widget to a page from the menu *Add* (➕) → *Widgets* → *Collaboration*. Click the button *Invite members to this site* to bring up the interface for sending invitations.

Click the plus sign next to a User or click the *Add Email Address* button to add a User to the invite list. Users that have already been invited but have not yet responded appear with a check mark next to their names. You can also invite Users to the *Site Owner*, *Site Content Reviewer*, and *Site Administrator* Roles for your site by selecting that Role under the *Invite to Role* heading. Once you've added all the Users you want to invite and have selected their Roles, click the *Send Invitations* button to invite them. For more information on roles, see the Roles and Permissions documentation.

The Site invitation shows up under the *Requests List* tab on the User's *Notifications* page. The User can then choose to *Confirm* or *Ignore* the invitation.

When Users confirm such invitations, they become Site members assigned to the Roles you defined.

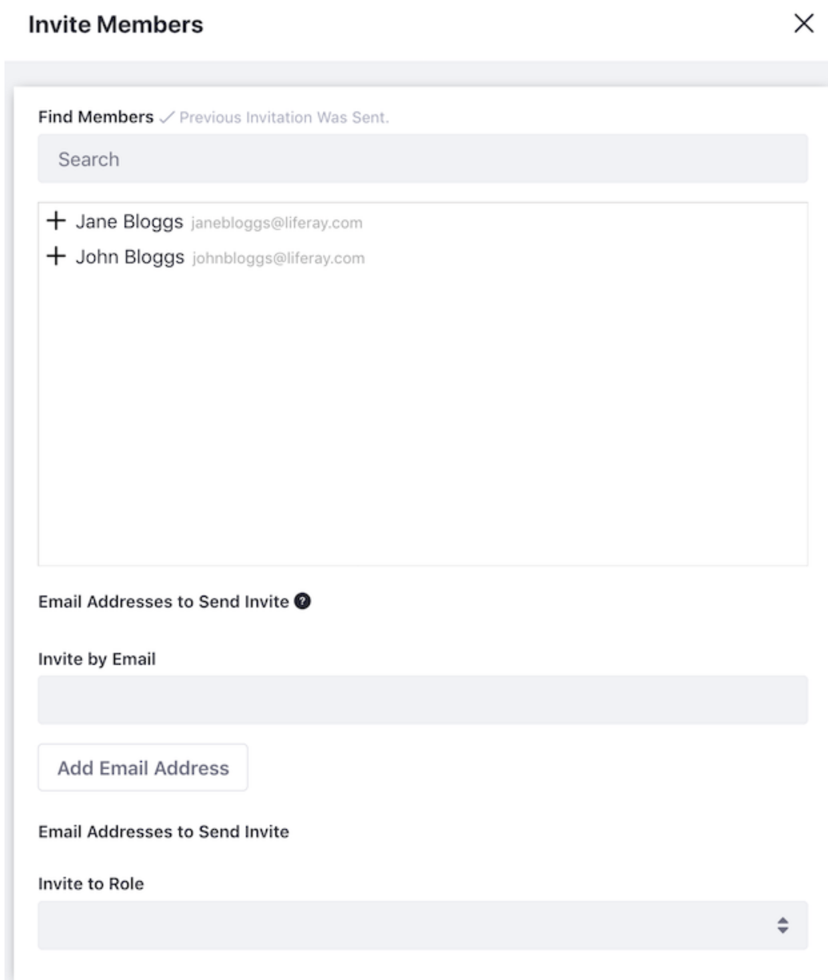

Figure 54.21: You can invite users by clicking the add sign next to the user's name.

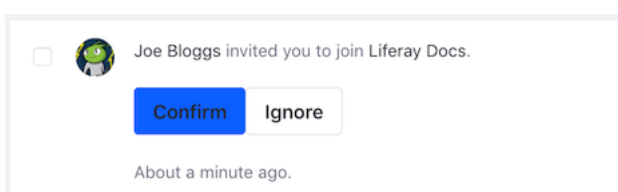

Figure 54.22: You can confirm or ignore the invitation.

Chapter 55

Creating A Social Network

Liferay DXP contains several features and widgets for leveraging its social framework. The Activities widget lets you broadcast user activities on a Site. This is a good way for Site members to see what's going on in their communities. When placed on a user's private Dashboard, the Activities widget displays the activities of that user's social connections. Users can make those social connections via the Contacts Center app, which lets them establish connections and followers throughout the portal. What's more, widgets can be exported as OpenSocial gadgets and/or used with Facebook.

This guide shows you how to do these things, and more.

55.1 Using the Activities Widget

The core social widget is Activities. It displays information about user activity on the Site where you added it. User activities tracked include updates to the Documents and Media library, blog posts, message boards posts, wiki pages, and bookmarks. Liferay DXP also tracks information about web content but only displays this information if the logged-in user is a Site administrator. This widget provides a summary of recent Site activity. You can use it on a Site's public or private pages to show what Site members have been up to, or you can use it on the public or private pages of a user's personal Site. When added to a personal Site, the Activities widget shows the activities of only that user.

Add the Activities widget to a page from the *Add* ( ) → *Widgets* → *Social* menu.

Note that the widget provides links to the assets listed in the feed. These links don't work, however, unless there's a way to display the assets on the page. For example, if you click a link to a blog post in the Activities widget, that page must have a Blogs widget to display that blog post.

55.2 Connecting Users

By adding the Contacts Center and Activities widgets to users' private Dashboards, administrators can let users form social connections. The Contacts Center widget lets users form these connection types:

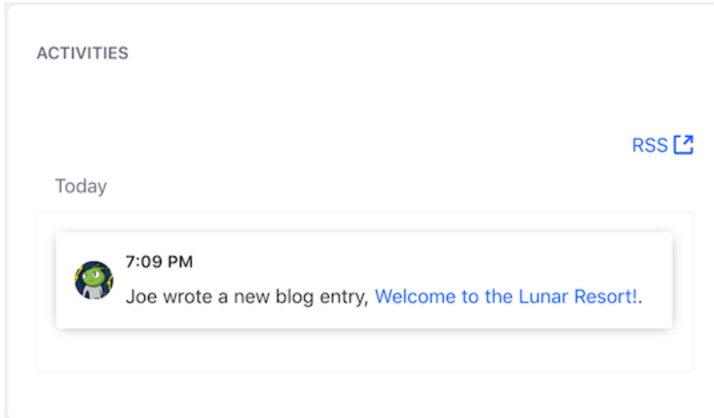

Figure 55.1: The Activities widget shows information about asset-related user activity in the current Site.

- **Connection:** A two-way relationship formed by a user accepting a connection request from another user. The Activities widget on each user's private Dashboard displays the activities of the other user.
- **Following:** A one-way relationship in which one user follows another user. The followed user's activities appear in the Activities widget on the follower's private Dashboard.

To ensure that all users have a page on their private Dashboard that contains the Contacts Center and Activities widgets, you can create a user group and then create a page via a Site Template. For instructions on this, see the user groups documentation. The widgets are in *Add* (+) → *Widgets* → *Social*.

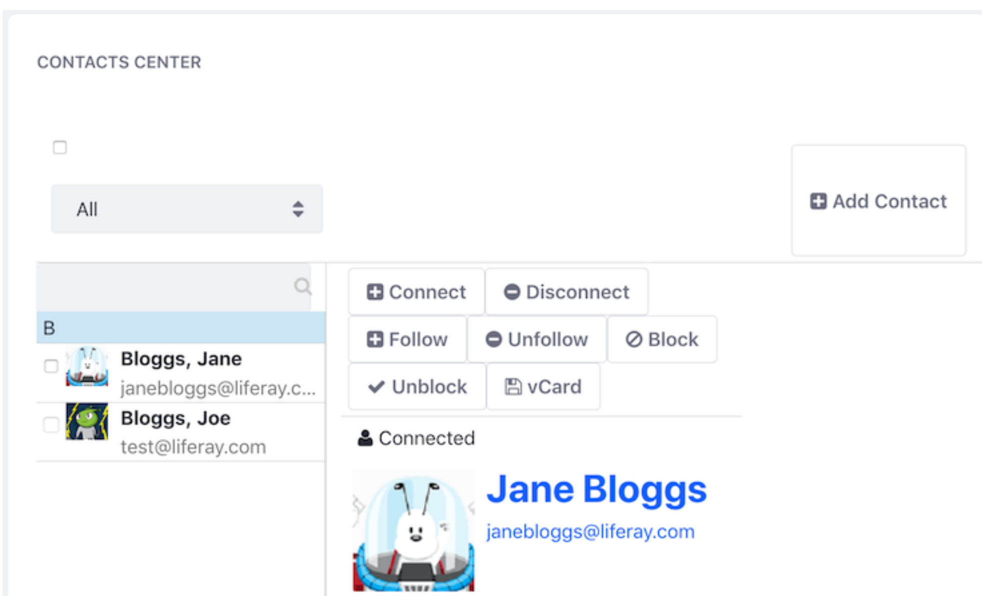

Figure 55.2: The Contacts Center widget lets users make connections.

When you select a user in the Contacts Center, these buttons appear across the top of the widget:
Connect: Send a connection request to the user.

Disconnect: Disconnect from the user.

Follow: Follow the user.

Unfollow: Stop following the user.

Block: Block the user. Blocking a user only prevents that user from following you or adding you as a connection. A blocked user can still send messages to and view the public profile information of the blocking user.

Unblock: Stop blocking the user.

vCard: Export the user's vCard and save it as a VCF file. vCard is a file format standard for electronic business cards.

When you send a connection request, the user is notified and can confirm or ignore the request.

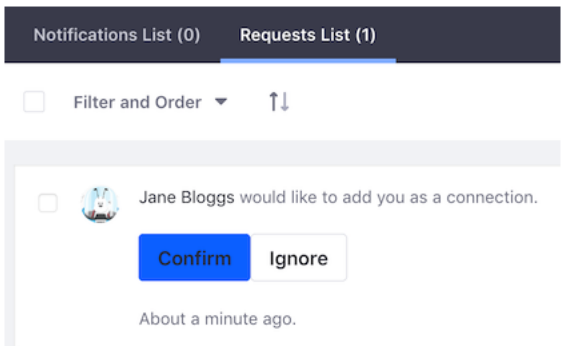

Figure 55.3: Users get a notification that lets them respond to connection requests.

55.3 Exporting Widgets To Other Websites

You can publish its widgets to other websites via OpenSocial. This lets you provide your widget or content in the context of another website. Read on to find out how this is done.

Sharing OpenSocial Gadgets

OpenSocial consists of a set of APIs for social networking. It may be beneficial for you to share widgets from your server with other sites, such as ighome.com or igoogleportal.com. These sites let users customize their own pages with gadgets. Users can share Liferay DXP widgets on any OpenSocial-compatible site. Liferay DXP does this by sharing the widget's URL with the OpenSocial platform. This URL is unique to the user's specific widget instance. Users can therefore share multiple instances of the same widget as different OpenSocial gadgets.

Follow these steps to share a widget with OpenSocial:

1. Add the widget you want to share to a site page.

2. Click the widget's *Options* icon () and select *Configuration*.

3. In the *Sharing* tab, open the *OpenSocial Gadget* section.

4. Set the selector to *YES* for *Allow users to add [widget name] to an OpenSocial platform*. In the *OpenSocial Gadget URL* field, replace `localhost:8080` with the name of your public domain and port.

5. Click *Save*.

6. Close the dialog and click the widget's *Options* icon ( ). There's a new option available named *Add to an OpenSocial Platform*. Select this option to add the widget to a page on an OpenSocial platform.

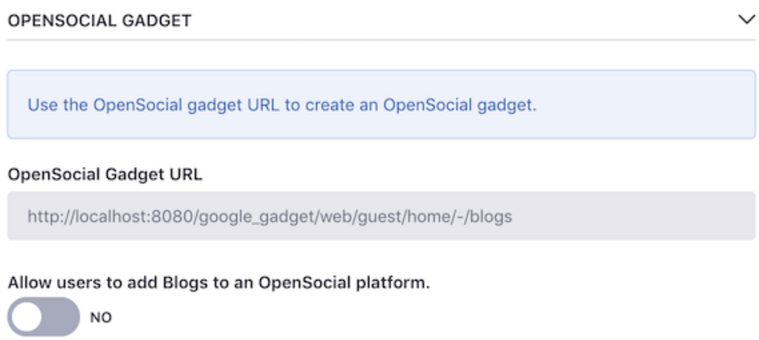

Figure 55.4: You can share widgets via OpenSocial.

55.4 Integrating with Facebook

Liferay DXP provides tools for integrating your portal and its content with Facebook. For example, you can use Facebook for authentication and even export widgets as Facebook applications. This article shows you how.

Facebook Sign On

Like many websites you may visit, any site running on Liferay DXP can use Facebook for sign in. This makes it easier for users to sign in to your Site, since they don't need to remember another user name and password. See the deployment documentation for information on using Facebook to sign into Liferay DXP.

Using Your Widget as Facebook Applications

You can add any Liferay DXP widget as an app on Facebook. To do this, you must first get a developer key. A link for doing this is provided to you in the Facebook tab in any widget's Configuration screen. You must create the app on Facebook and get the key and canvas page URL from Facebook. You can then copy and paste their values into the Facebook tab. Once you do that, your widget is available on Facebook.

This integration lets you make things like Message Boards, Calendars, Wikis, and other content on your website available to a much larger audience (unless you already have a billion users on your site, in which case, kudos to you). If you're a developer, you can implement your widget on Liferay DXP and then publish it to Facebook.

55.5 Using Social Bookmarks

Social bookmarks appear below content as buttons for sharing that content on social networks. For example, social bookmarks appear by default in the Blogs widget below each blog post. For more information on configuring social bookmarks in the Blogs widget, see the documentation on displaying blogs.

These social bookmarks are available by default:

- Twitter
- Facebook
- LinkedIn
- Google Plus

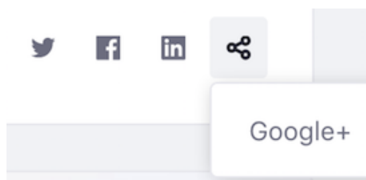

Figure 55.5: The default social bookmarks appear inline below content.

You can install the Social Bookmarks app from Liferay Marketplace. This app is available for Liferay CE Portal and Liferay DXP. It adds the following social bookmarks:

- AddThis
- Delicious
- Digg
- Evernote
- Reddit
- Slashdot

If you need help installing apps from Liferay Marketplace, see the article on using Marketplace.

Chapter 56

Search

Sites often feature lots of content split over lots of asset types. Web content articles, documents and media files, and blogs entries are just a few examples. Most content types are *assets*. Under the hood, assets use the Asset API and have an Indexer class. Any content that has these features can be searched.

Figure 56.1: The Type Facet configuration lists the searchable out-of-the-box asset types.

56.1 Elasticsearch

The default search engine is Elasticsearch, which is backed by the Lucene search library. There's an Elasticsearch server embedded in all bundles, which is handy for testing and development purposes. Production environments must install a separate, remote Elasticsearch server (or even better, cluster of servers). For information on how to install Elasticsearch, read the deployment guide.

56.2 Search Features

Searching is simple and straightforward. Find a search bar (there's one embedded in every page by default), enter a term, and click *Enter*.

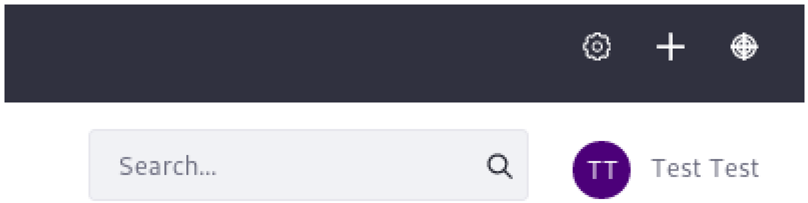

Figure 56.2: There's a search bar embedded on all pages by default.

Search Modularization: If you're upgrading from an earlier Liferay DXP version, you'll expect the embedded search bar to be the entry into the Search application. However, the search functionality is modularized now, and the Search Bar is a standalone application. See the article on new search features for more information.

After search is triggered, a results page appears. If there are hits to search engine documents, they appear as search results in the right hand column. In the left hand column are search facets.

The search bar, search results, and search facets make up three powerful features in the search UI.

Search Bar

The search bar is simple: it's where you enter *search terms*. Search terms are the text you send to the search engine to match against the documents in the index.

Search Results and Relevance

The search term is processed by an algorithm in the search engine, and search results are returned to users in order of relevance. Relevance is determined by a document's *score*, generated against the search query. The higher the score, the more relevant a document is considered. The particular relevance algorithm used is dependent on algorithms provided by the search engine (Elasticsearch by default).

Search Facets

Facets allow users of the Search application to filter search results. Think of facets as buckets that hold similar search results. You might want to see the results in all the buckets, but after scanning the results, you might decide that the results of just one bucket better represent what you want. So what facets are included out of the box?

- Category
- Folder
- Site

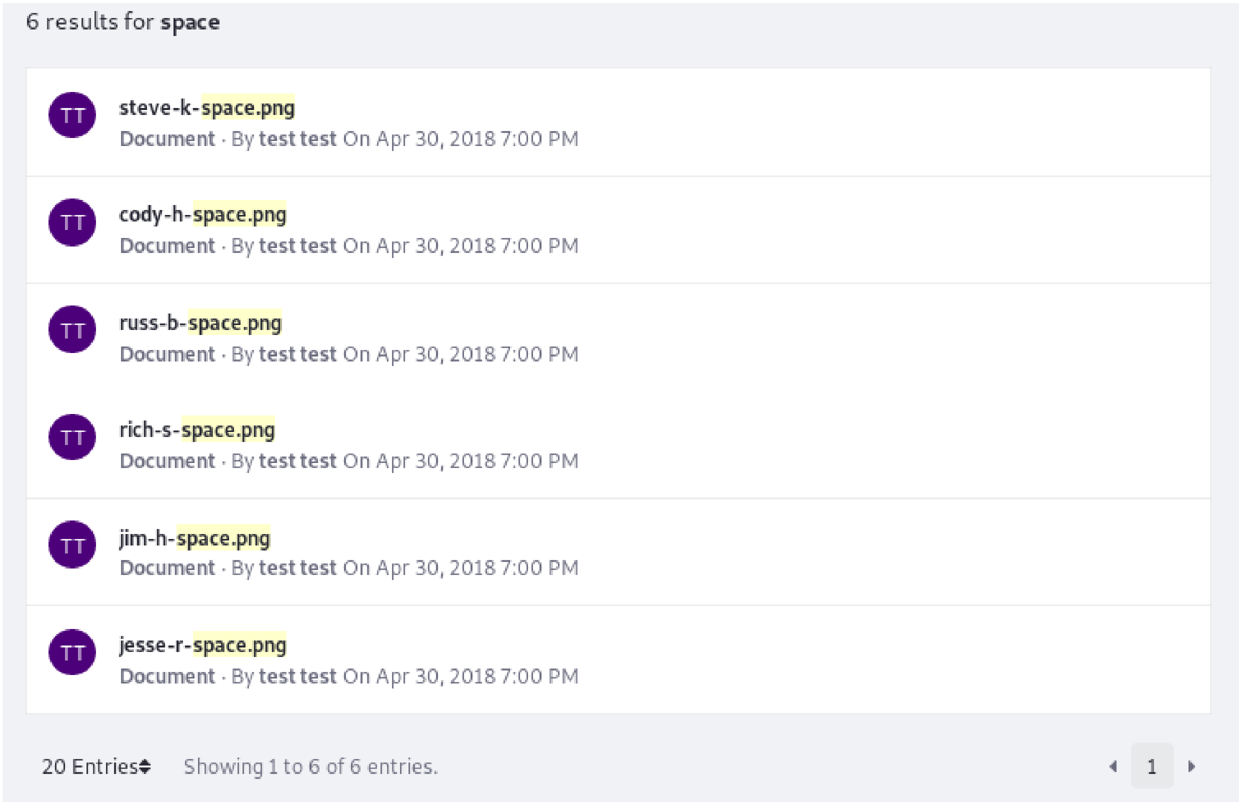

Figure 56.3: Results are displayed in the Search Results portlet.

- Tag
- Type
- User
- Modified
- Custom

You've probably used something similar on any number of sites. You search for an item, are presented with a list of results, and a list of buckets you can click to further drill down into the search results, without entering additional search terms. Search facets work the same way. Facets are, of course, configurable.

56.3 What's New with Search?

There are lots of new features in the search functionality. This brief overview can get you familiar with the changes. More details are in the remaining articles.

Modularized Search Functionality

In prior versions, all the search functionality was in one place: the Search portlet. Now each separate feature is available in a separate portlet, and you can drag and drop search widgets onto a

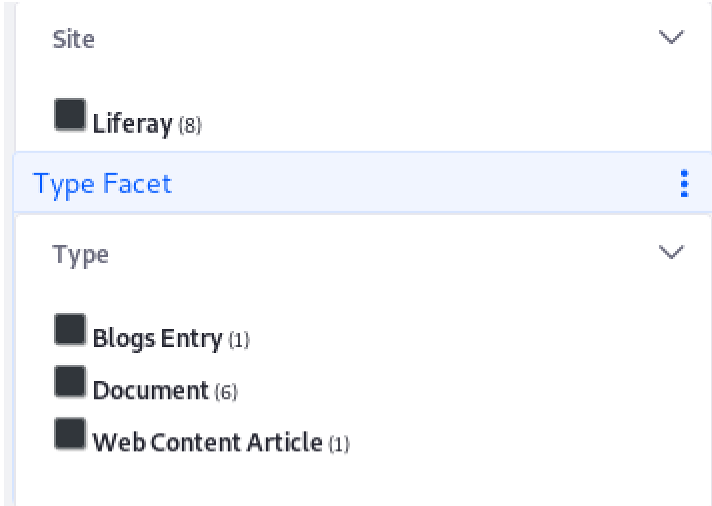

Figure 56.4: *Site* and *Type* are two of the facet sets you'll encounter. They let you drill down to results that contain the search terms you entered.

page, creating a perfectly customized search page. The search configurations that existed in the previous version are still available, but the setting you're looking for might be in a slightly different location due to the modularization of search in 7.0. The articles here help familiarize you with the new look.

Elasticsearch 6

Support for Elasticsearch 6 was added as an *opt-in* feature for Liferay DXP version 7.0. In the current version, Elasticsearch 6 is the default, embedded search engine. See the article on Installing Elasticsearch to learn how to install a remote Elasticsearch server for a production environment.

Search Administration

Carry out search administration tasks (like reindexing the search indexes) in the Search administration application in the Control Panel (Control Panel → Configuration → Search).

Default/Customizable Search Pages

Use a pre-configured search page template to get a sensible default search page up and running.

Advanced Search Syntax

Use Elasticsearch's Query String syntax to create advanced search queries (if using Elasticsearch as your Search Engine). The actual query syntax and further processing are dependent on the search engine's implementation details. Consult your search provider's documentation for more information.

Multiple Facet Term Selection

Narrow search results using Facets. Continue refining results by selecting as many facets terms as necessary. Previously, only one facet term could be used at a time per each facet.

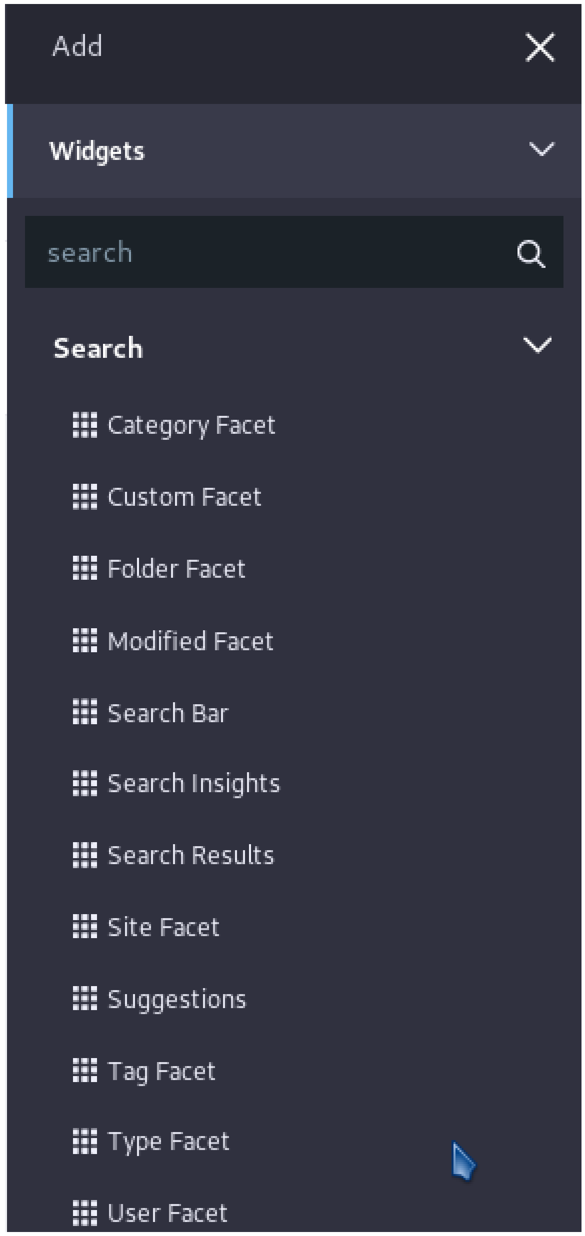

Figure 56.5: The search functionality is now distributed across several widgets.

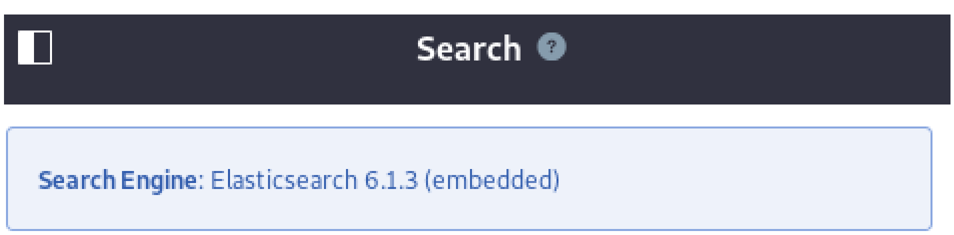

Figure 56.6: Elasticsearch 6 is the default search engine.

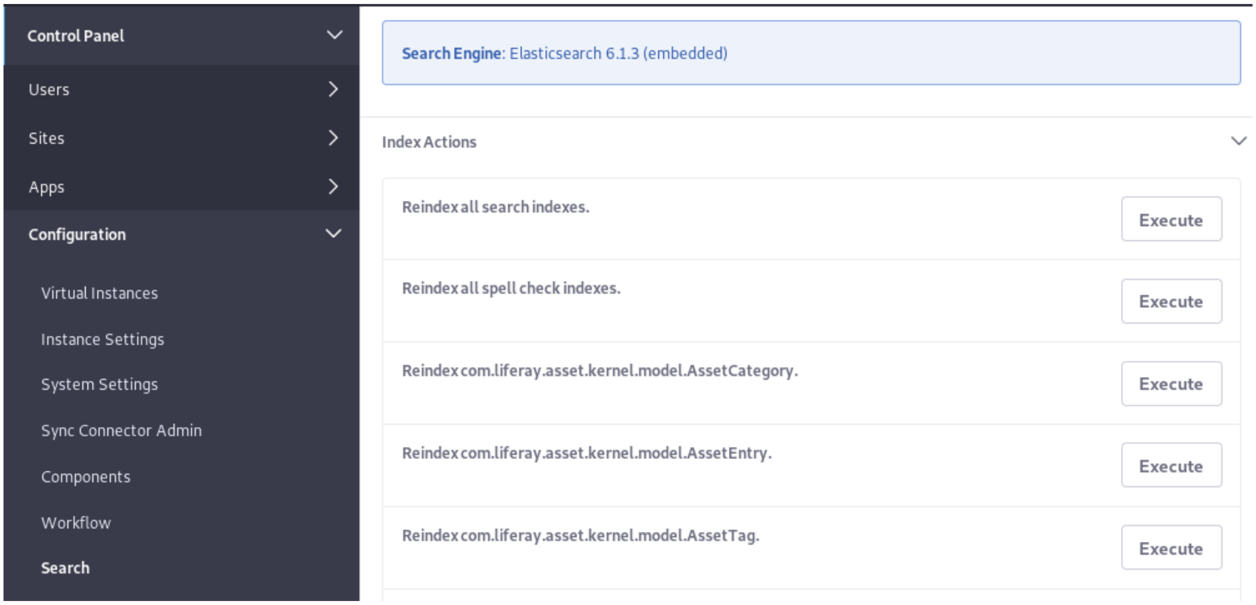

Figure 56.7: Reindexing content now happens in the Search administration application.

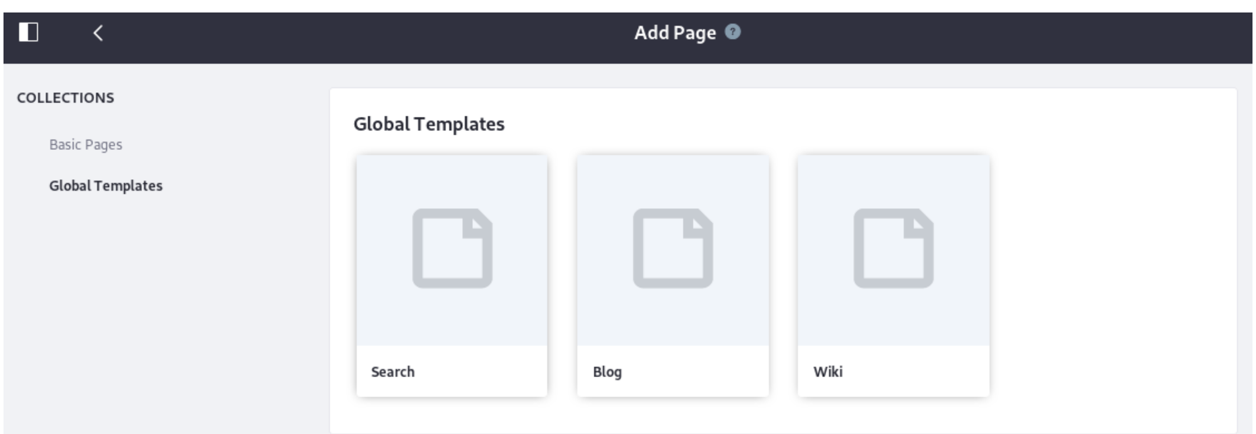

Figure 56.8: Use the search page template to create your site's dedicated search page.

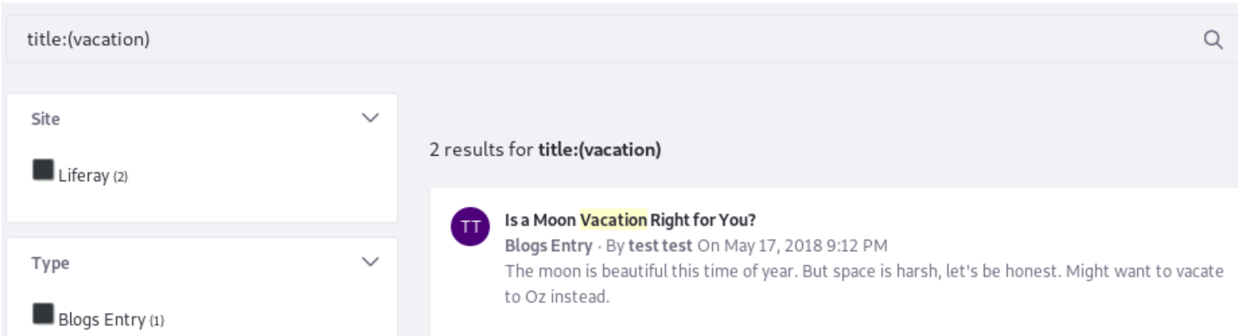

Figure 56.9: Use Elastic's Query String syntax to construct advanced search queries.

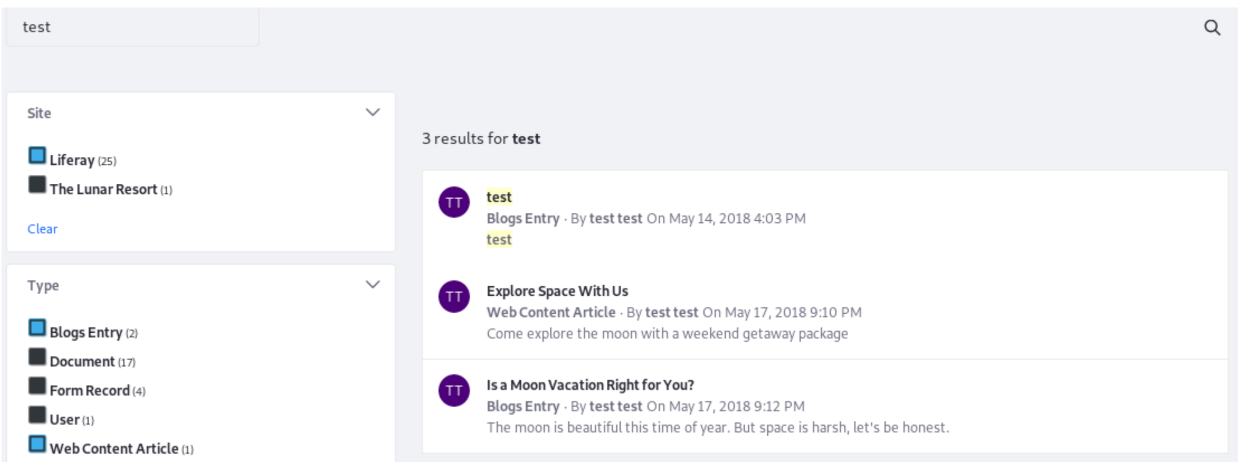

Figure 56.10: Facets and their terms can be selected in multiples to refine search results.

Search System Settings

Configure system wide search behavior from the Search System Settings category found in Control Panel → Configuration → System Settings → Search. The Search Web entry reverts the default search experience from using the new Search Widgets to the classic Search Portlet that was standard in past releases.

Custom Facets

Configure a custom facet to create a new result aggregation on any non-analyzed keyword field.

Search Options

Configure page scoped search behavior by adding the Search Options widget from the Add Widget menu.

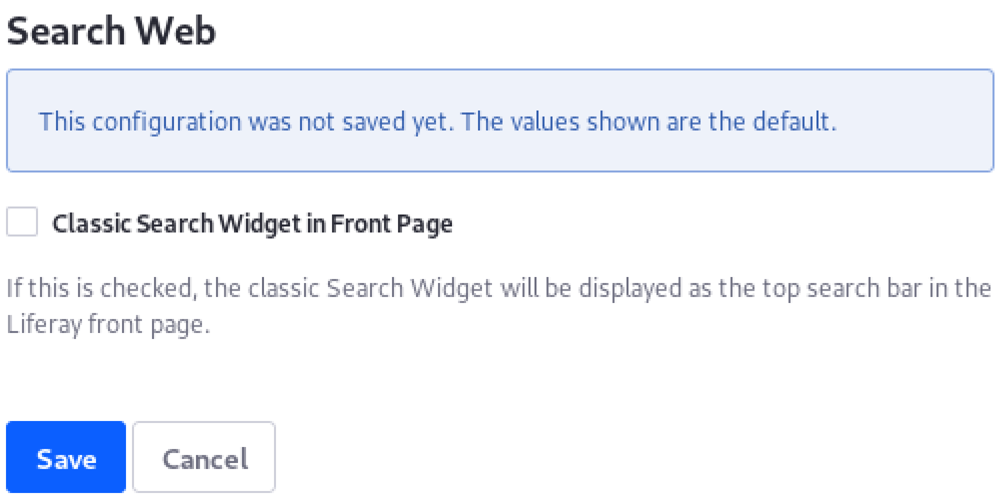

Figure 56.11: Use the Search Web entry in System Settings to enable the classic, monolithic Search widget.

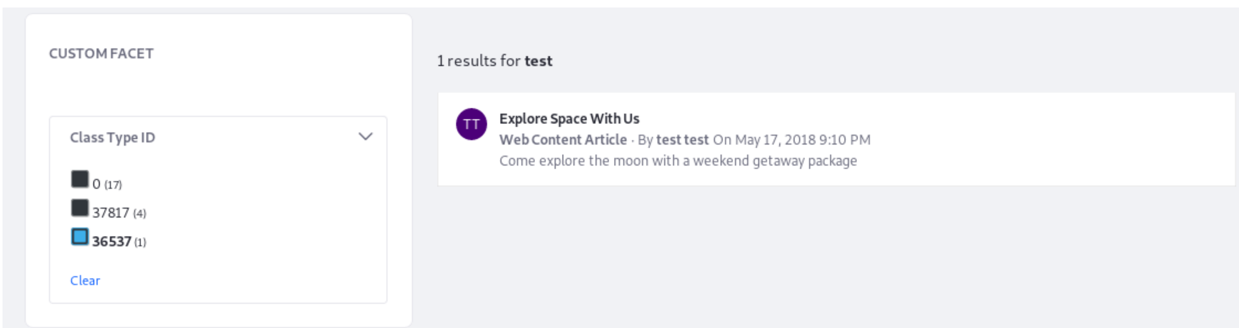

Figure 56.12: Configure a custom facet on the Class Type ID field.

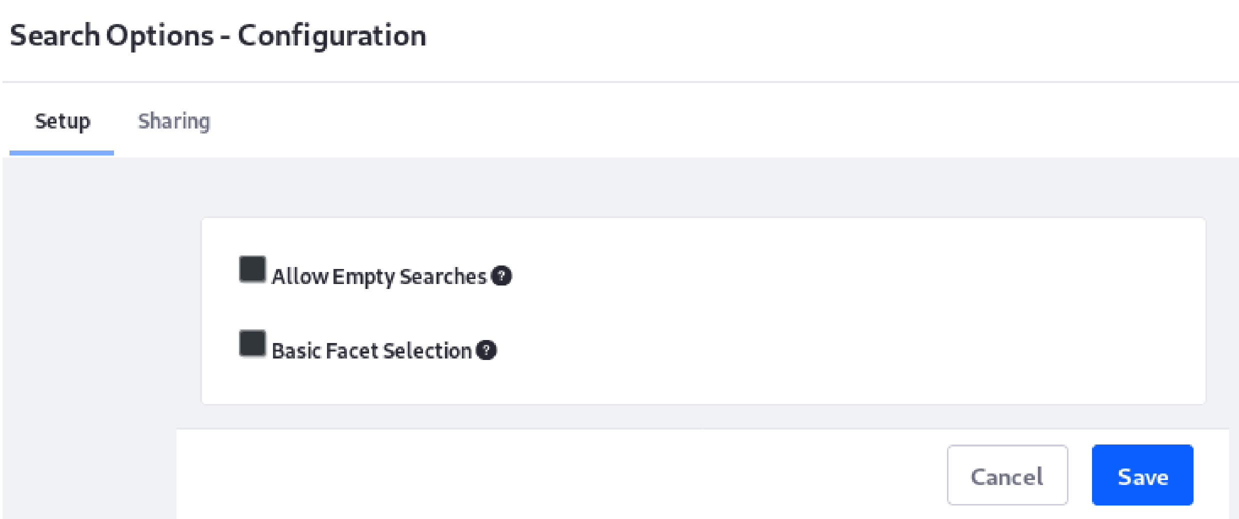

Figure 56.13: The Search Options widget configures the search experience for the page it's deployed on.

Search Facet Friendly URLs

Using the new Search widgets, filter some search results by applying a facet term. The URL is now friendly:

```
http://localhost:8080/web/guest/search?q=test&modified=past-hour
```

Code Changes and Improvements

See the developer guide to learn about the new, deprecated, and removed APIs and extension points for search infrastructure.

Solr 7.x Support

The Solr Connector application was updated to support Solr 7.x.

Read here for more information.

Legacy Search

We're cheating with this one. It isn't a new feature, but the old way of doing search, inside one standalone application, is still just one checkbox away (located in Control Panel → Configuration → System Settings → Search → Search Web). This causes the embedded search application to revert to the old-style search portlet. The legacy search application can also be added to any page from the Add Widget menu (it's under the *Tools* category).

To learn more about using and configuring these new features, see the remaining documentation in this section.

56.4 Configuring Search Pages

There are multiple ways to skin the search cat (disclaimer: no actual cats have been harmed during the writing of this article).

If you're unsure which approach to take, use the default configuration. It provides a sensible starting point that can be modified later, as needed.

If you've been using Liferay DXP for a long time and like the search experience you've always used, use the legacy approach.

If you're in need of a fully customized experience, manually configure the search experience.

After choosing your approach and reading here to get it up and running, find the articles on the Search Bar, Search Facets, and Search Results to understand the full suite of configuration options.

Search display pages are where users go to enter search terms and browse search results.

Search Page Templates

The default search page is backed by a Search Page Template, and manually configured search pages can use the template, too. The template can be used in two ways:

1. Enable inheriting changes to the template, if you want the search page to get any updates made to the template at a later date.

2. Create the page based on the template, but independently configured after the initial creation.

Out of the box, the Search Page Template includes these widgets:

- Search Bar
- Suggestions
- Search Results
- Search Options
- Site Facet
- Type Facet
- Tag Facet
- Category Facet
- Folder Facet
- User Facet
- Modified Facet

Out of the box, widgets use the *Barebone* Application Decorators: unless there's content to render in the widget, the widget body is hidden. The header is displayed if you hover over it.

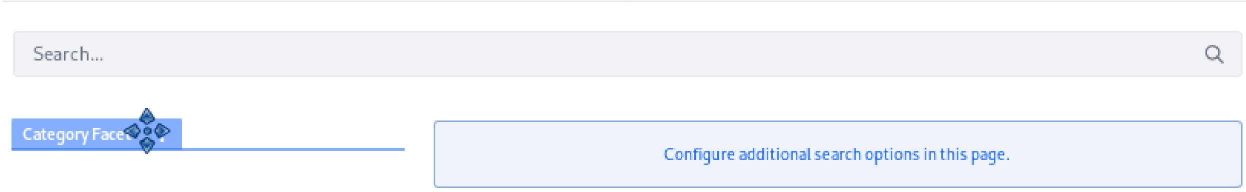

Figure 56.14: At first glance, not much is happening on the search page. But, there's more than meets the eye.

Because of this, when you visit a search page created from the default search page template, you won't see certain widgets fully rendered.

By contrast, when you add a search widget to a page manually, they use the *Borderless* decorator (by default), which shows more of the widget even when there is no content to display.

Default Search Pages

Using the default site and the default theme with the default search settings, the out-of-the-box search experience has two components for end users:

1. A search bar embedded on each page.
2. A default search page where search requests are routed.

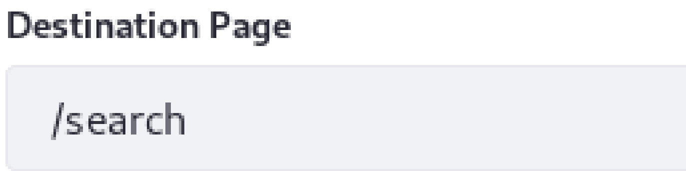

Figure 56.15: By default, the embedded search bar points to the pre-configured /search destination page.

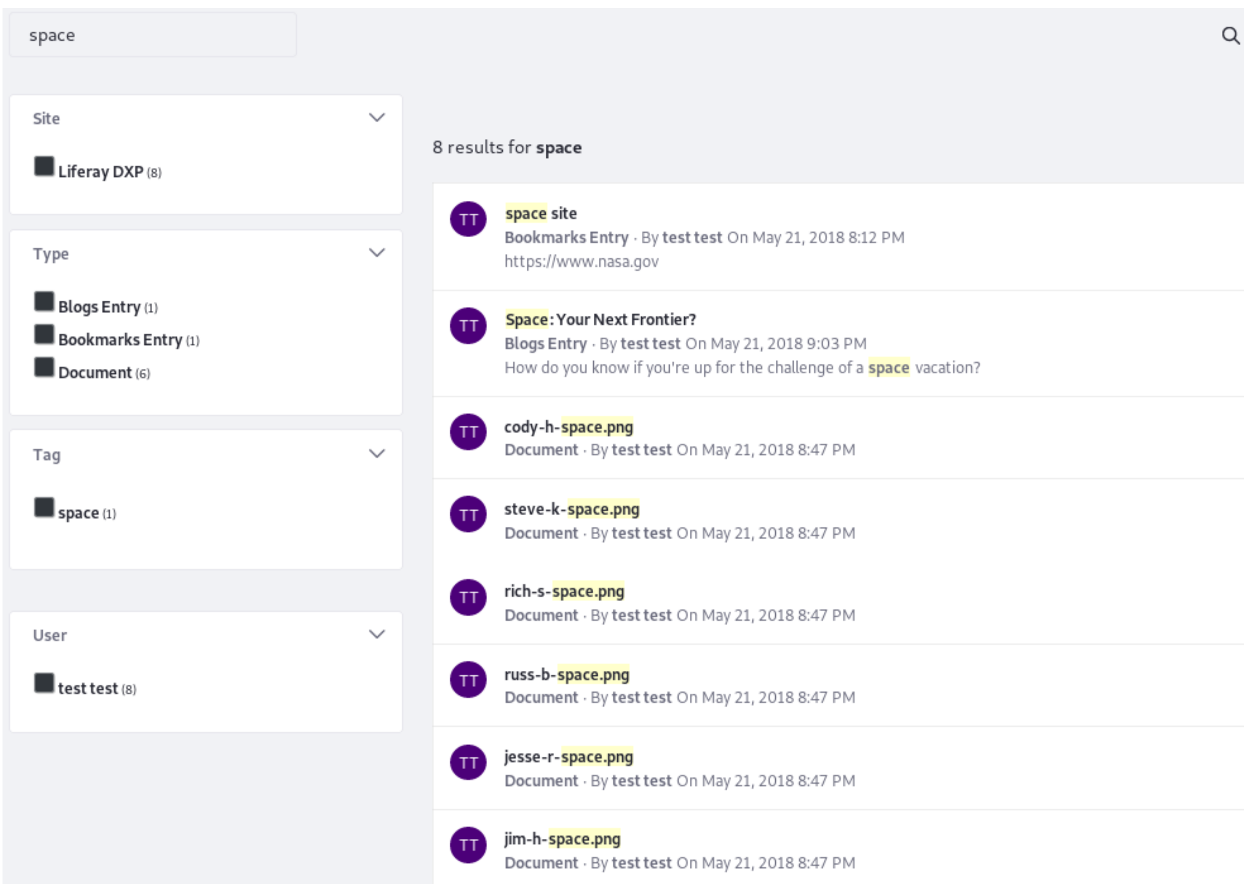

Figure 56.16: The default page is pre-configured with the Search Results widget and the various Facet widgets to provide a full search experience.

Behind the scenes, The search bar widget points to a hidden search page with the friendly URL /search.

Enter a search term and you're redirected to the default search page, where results are displayed in the Search Results widget.

The default search page is based on a Search page template, but it doesn't inherit changes from the page template by default. That means you can customize the search page directly without changing the template's inheritance configuration.

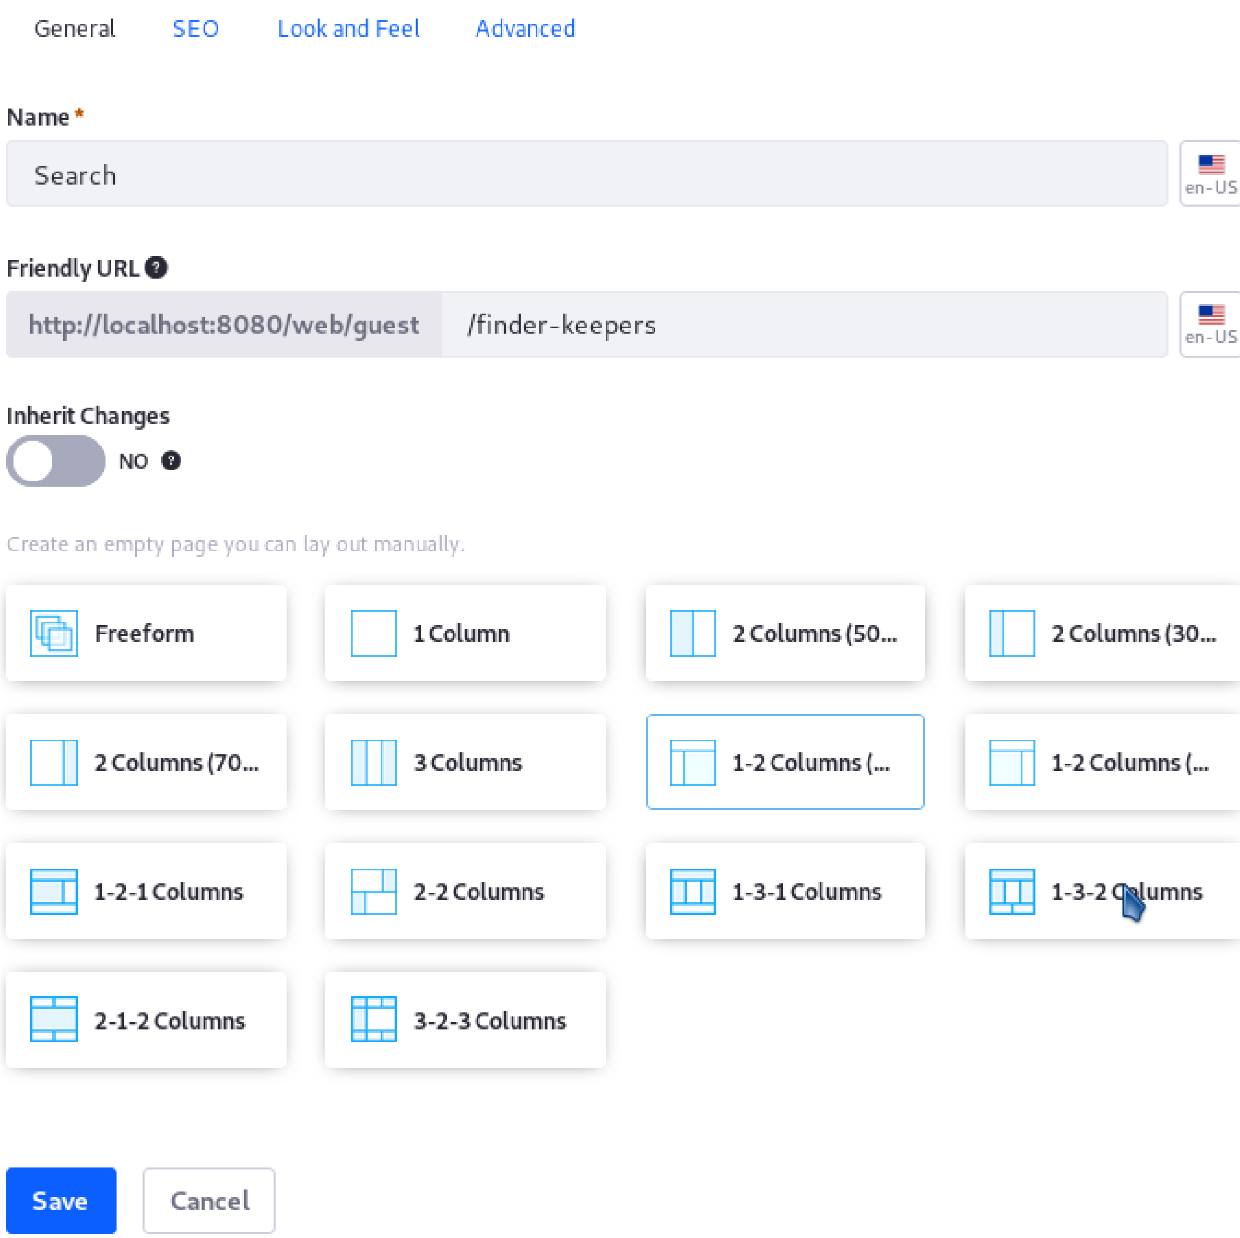

Figure 56.17: Configure the Search page. By default, it doesn't inherit changes from the page template.

If you require just a few changes to the default page, don't abandon it and create one manually. Just make the configuration changes you need, including adding, configuring, and removing widgets on the page. On the other hand, if you want a clean break from the default search page, starting from scratch is also an option.

Manual Search Page Configuration

It's reasonable to create the search experience from the ground up. If you're working from a newly created site, it's a necessity. These steps show you how to switch to a manually configured search experience in the default site, but you can skip the step on deleting the default search page if you're starting with a new site:

1. Delete the existent search page by navigating to the default site's menu and clicking *Build* → *Site Pages*. Click the Search page's Options menu icon ( ) and select *Delete*. Confirm you want to delete the page, and it's gone.

   Once deleted, the search bar disappears from your site pages, replaced by a warning message visible only to site administrators:

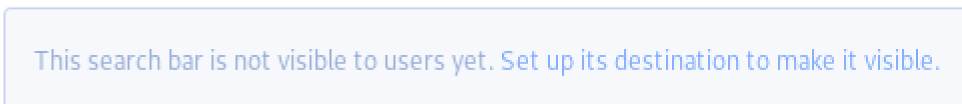

Figure 56.18: The search bar is only visible if it points to an existent page.

2. Create a new page named whatever you want (*Finders Keepers*, perhaps). Make it hidden or add it to the navigation as you please (the default search page is hidden from the navigation).

   If you want a pre-configured search page, create it from the Search page template. Find the template in the Add Site Page form. It's under *Global Templates*.

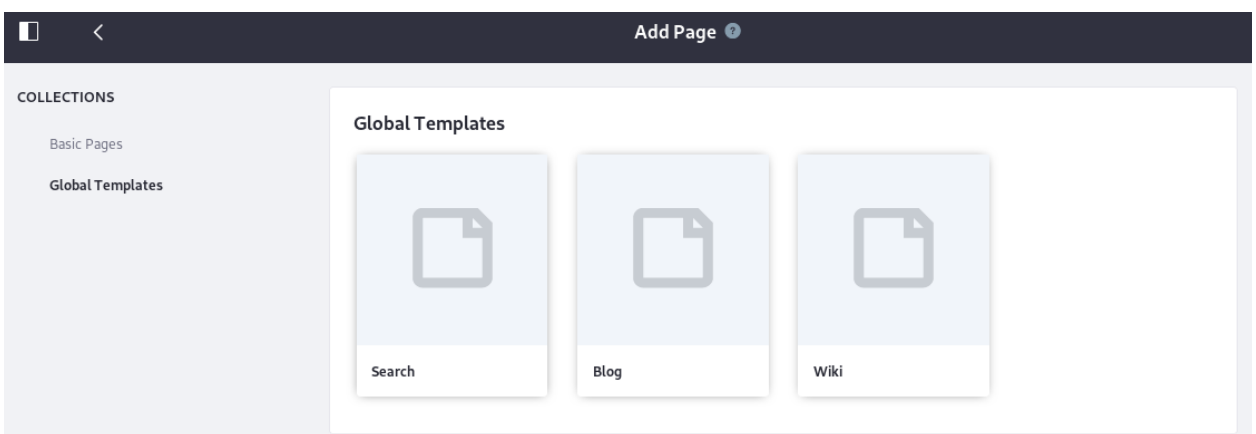

Figure 56.19: There's a handy page template for creating search pages.

3. If you're creating a page not backed by the template, add and configure all the widgets you need. You'll find all the available search widgets in the Add Widget menu's Search section. Lay them out however you want on the page.

4. Configure the search bar at the top of the page, making sure it points to your new search page's friendly URL (for example, `/finders-keepers`).

   Click the Search Bar widget's Options menu (⋮).

   Click *Configuration* and set the Destination Page to the search page's friendly URL.

   Click *Save*.

Now your search page is up and running.

Legacy Search Experience

In prior versions, the search experience was encapsulated in one application, *Search*. It was embedded in the default theme, just like the search bar is now. It looked very similar, with only the search bar visible in the default view of the application. Once a search term is entered, the maximized view of the application is presented, with all the search facets and results now in view. It looks a lot like the new search behavior, only its monolithic structure means it's difficult to customize. If you liked the old application, it's still available. Enable it with these steps:

1. Delete the default search page. From the site menu, click *Build* → *Pages*. Click the Options menu (⋮) for the Search page and choose *Delete*.

2. Enable the legacy search application. Go to Control Panel → Configuration → System Settings → Search → Search Web and check the box for *Classic Search Widget in Front Page*.

Now your portal's search is backed by the legacy Search application, and it's embedded on each page in the default theme. To add the legacy Search application to a page, open the Add Widget menu, find the Search widget under the Tools category, and drag and drop it onto the page.

Configure the portal's search behavior to suit your needs. Here you've seen three distinct search configurations.

56.5 Searching for Assets

As explained in the Search introduction, all indexed assets can be returned as search results. Developers can create their own assets, so your installation might have additional asset types beyond the ones included by default.

Searching for Users: When you click an asset in the search results, it's displayed in an Asset Publisher (unless the *View in Context* option is selected in the Search Results portlet). Users are different, though. Think of them as invisible assets, not intended for display in the Asset Publisher application. While Users appear as search results with other indexed assets, when you click one you're taken to the User's profile page. If public personal pages are disabled, clicking on a User from the list of search results shows you a blank page.

Search Bar

Users enter the search context in the search bar. Users enter search terms, hit their *Enter* button (or click the magnifying glass icon), and they're taken to a search page with various search widgets deployed.

If using the Search Bar in the legacy search portlet, users see a maximized view of the search portlet displaying any results and facets that apply. See the article on configuring search pages to learn more about these options.

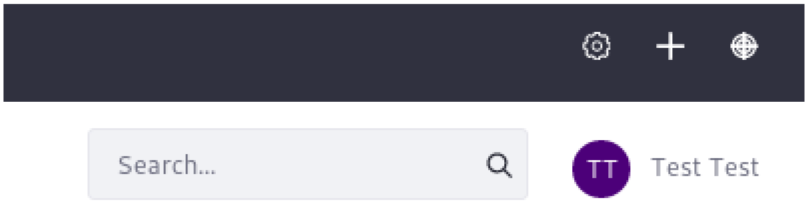

Figure 56.20: The default search configuration displays a search bar in its default view, beckoning users to enter the search context.

Entering Search Terms

Liferay's search infrastructure supports full text search as implemented by its supported search engines (Elasticsearch and Solr).

Full text search compares all the words entered in a search query (for example, *space vacation*) to all the words in each index document. A search engine like Elasticsearch calculates relevance scores to ensure the best results are returned first (like a Blogs Entry titled *Is a vacation in space right for you?*) and lots of matching results are returned (anything with either the word *vacation* or *space* is returned).

In addition to full text search, advanced search syntax is supported. Liferay DXP relies on the underlying search engine for this behavior, so consult the Elasticsearch or Solr documentation for the details.

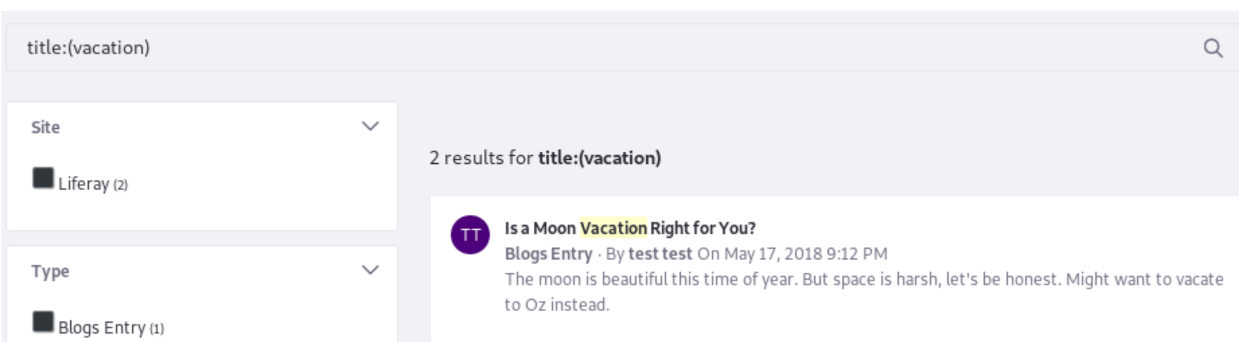

Figure 56.21: Search for text in a specific field using Elasticsearch's Query String syntax.

Matching Exact Phrases: Quoted Search

What if users want their search terms (for example, *space vacation*) to produce only results with the exact phrase, as typed? In a regular full text search, searching *space vacation* returns search results containing just the terms *space* and *vacation*, and hits containing both terms but separated by other text, as well as results with the exact phrase match. To ensure that only hits with the exact phrase are returned, enclose it in quotes: *"space vacation"*.

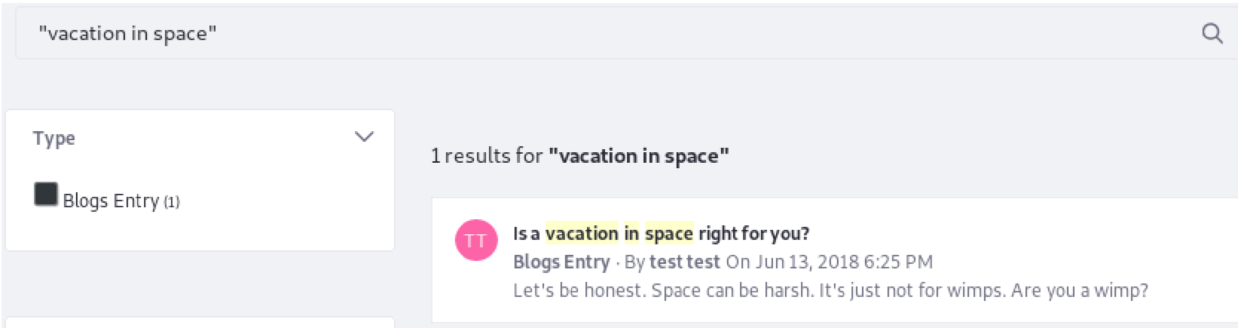

Figure 56.22: Search for exact phrase matches by enclosing search terms in quotes. If a user searched for *"space vacation"*, this result would not be returned.

Prefix Searching

If you're searching in a site for classical musicians, you might search for the term *instrument*. This search returns documents with the full word in them, but it also returns variants with *instrument* as the prefix. For example, results with *instruments*, *instrumental*, and *instrumentation* are also returned.

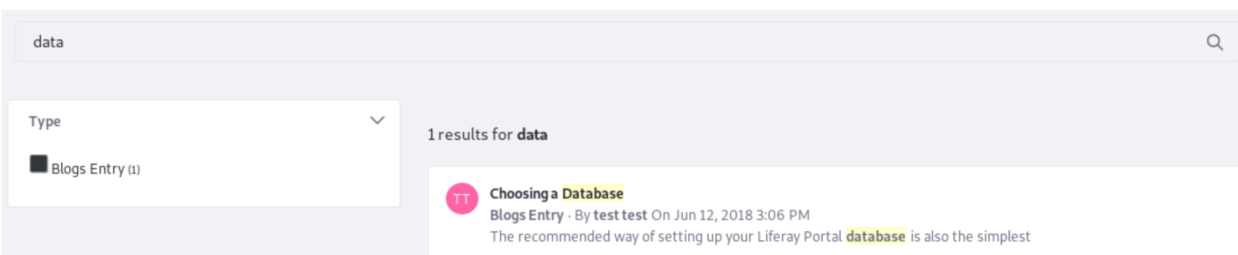

Figure 56.23: Searching for *data* also returns *database*.

Note: Prefix searching is available for many fields out of the box, but as with most things related to search behavior, it's more complicated under the hood. The details of the field mapping, including the analyzer used on the field and any transformations performed, determine the final behavior.

Another way to ensure users see results is through search suggestions.

Configuring the Search Bar

Configure the Search Bar's behavior via its portlet configuration screen.

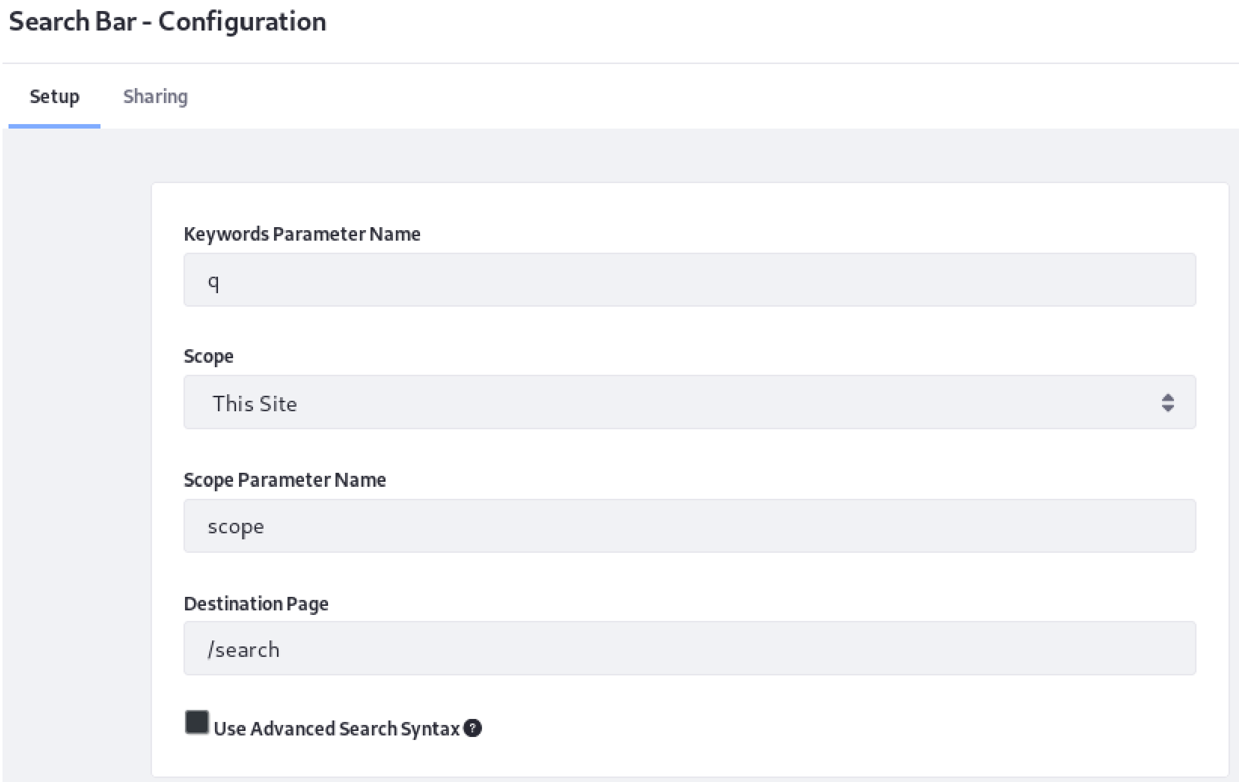

Figure 56.24: Configure the search bar behavior in its configuration screen.

Note: When you configure the globally embedded Search Bar widget at the top of one page, it configures the page-top Search Bar widget on all pages in the site. It also overrides the destination Search Page's Search Bar portlet, if they're configured differently. However, it does not override Search Bar widgets manually placed on other pages.

There are several options:

Keywords Parameter Name Edit the parameter name for the keywords entered in the search. For example, the default URL when searching for the keyword term *data* looks like this:

http://localhost:8080/web/guest/search?q=data

If you change the Keywords Parameter Name to *keyword* it looks like this:

```
http://localhost:8080/web/guest/search?keyword=data
```

Scope Choose between three options: This Site (default), Everything, and Let the User Choose. *This Site* means only the assets associated with the site where the search is executed are searched. Expand the scope of the search to all sites by selecting *Everything*. To let users choose which scope they want to search, select *Let the User Choose*.

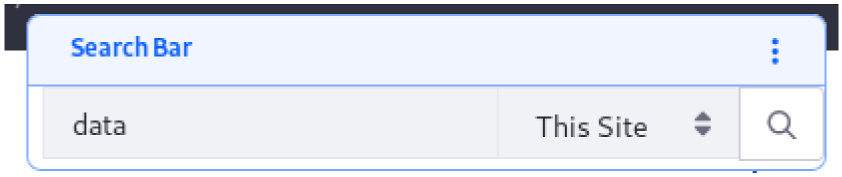

Figure 56.25: Let the user choose which scope the search is executed for.

Scope Parameter Name : Set the URL parameter name for the scope where the search is taking place. This parameter only appears in the URL if the scope *Let the User Choose* is selected. The default value is *scope*, so searching for the word *data* produces the default URL of

```
http://localhost:8080/web/guest/search?q=data&scope=this-site
```

Changing *scope* to *target* would produce this URL:

```
http://localhost:8080/web/guest/search?q=data&target=this-site
```

Destination Page Provide a friendly URL to the search page. If not configured or if it points to a page that doesn't exist, a message appears for administrators that the search bar must be configured for it to appear to users.

Use Advanced Search Syntax If using Elasticsearch, enabling this allows users to enter Query String Syntax into the Search Bar. If using Solr, consult its documentation for the proper syntax.

Search Suggestions

Suggest search terms to users when their initial queries are suboptimal. Spell check settings allow administrators to configure the Search application so that if a user types a search term that doesn't return many results (for example, a slightly misspelled werd), the user can be prompted to improve the search.

To configure the spell check settings,

1. You must first reindex the spell check indexes. Go to *Control Panel* → *Configuration* → *Search*, then click *Execute* next to *Reindex all spell check indexes*.

2. Add the Suggestions widget to the search page.

3. Open its configuration screen. Click the widget Options button () and select *Configuration*.

There are three main settings here:

Display "Did you mean..." if the number of search results does not meet the threshold. : Present users alternate, spell checked search queries if their search did not return a minimum number of results (50 by default).

Display Related Queries If the number of search results doesn't meet the specified threshold (50 by default), display up to a maximum number of alternative queries (10 by default).

Add New Related Queries Based on Successful Queries Index a user's search query if it produces a minimum number of results (50 by default), so it can be displayed to users as a suggestion. If the Display Related Queries setting is enabled, it's used as a related query for similar search queries that don't produce enough results.

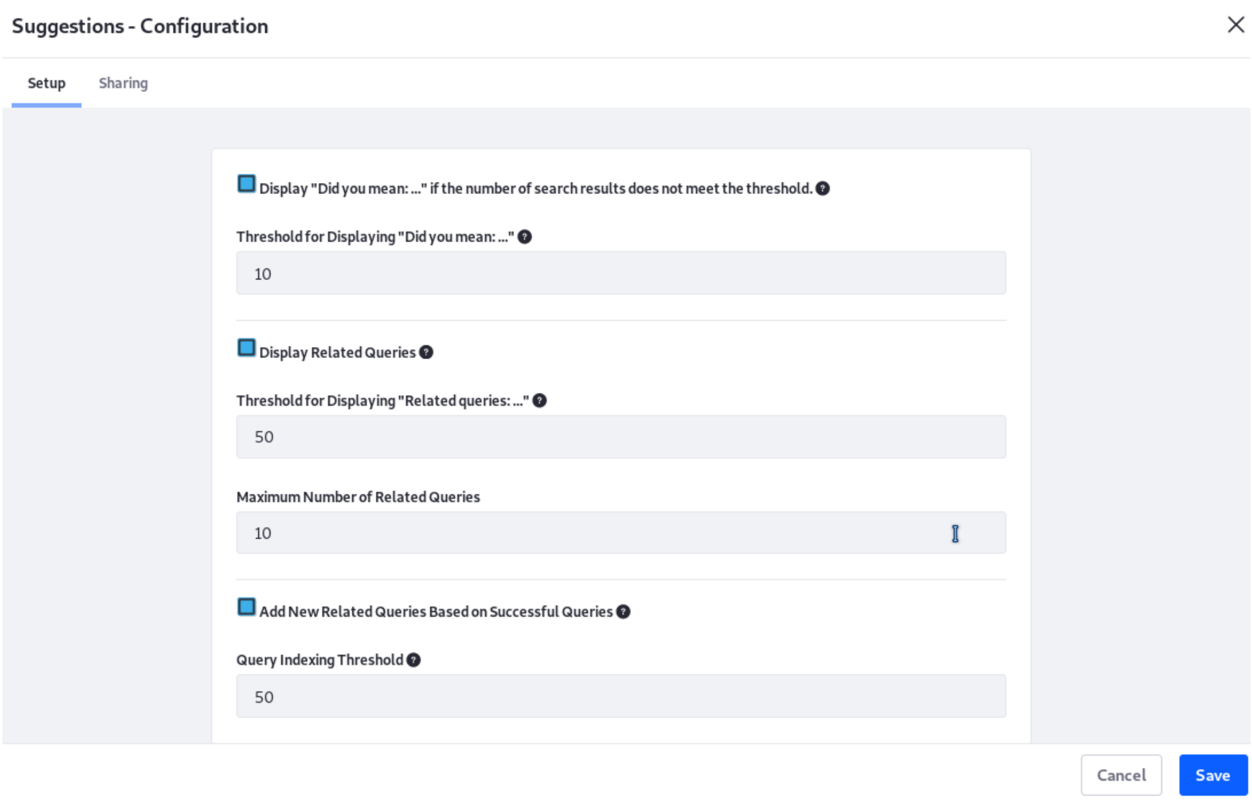

Figure 56.26: Configure the suggestion settings to allow for user input mistakes and help lead users to results.

CHAPTER 57

FACETS

Enter a keyword in the Search Bar and click the Search button (). The default search experience redirects to a page with results on the right and a collection of *facets* on the left.

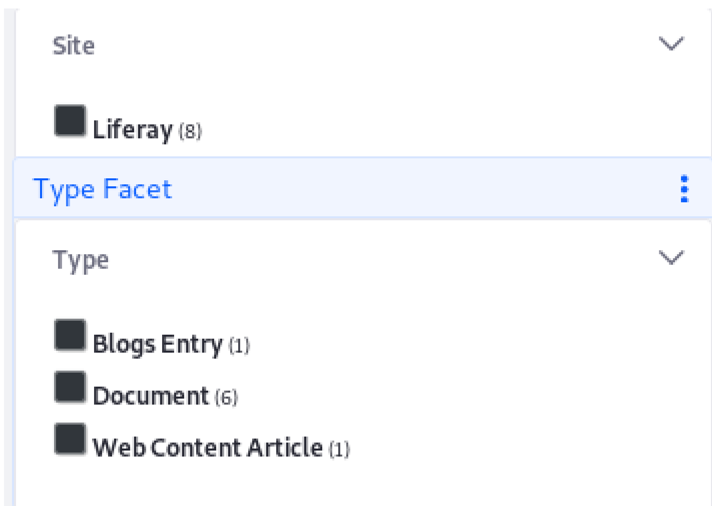

Figure 57.1: *Site* and *Type* are two of the facet sets you'll encounter.

A Facet aggregates search results by some common characteristic, with each facet holding search results that share something in common. After scanning the full list of results, a User might decide the results from just one facet are more appropriate (for example, all the results from a particular site, or all the results that are Blogs Entries). So what facets are included by default?

- **Site Facet** for filtering results by their site.
- **Type Facet** for filtering results by the Asset Type.
- **Tag Facet** for filtering results by Tag.
- **Category Facet** for filtering results by Category.
- **Folder Facet** for filtering results by Folder.
- **User Facet** for filtering results by the content creator.
- **Modified Facet** for filtering results by the Last Modified Date.

- **Custom Facet** for filtering results by some other indexed field. See here for more information.

Each item in a facet (selected using the checkbox) is called a *Facet Term* (*term* for short).

In this tutorial, you'll explore how facets and their terms are used and how to find a facet's configuration. The remaining articles show the configurations available for each facet.

57.1 Using Facets

If you're not actually an accomplished oboe player, pretend for a moment. You're visiting a site for classical musicians. You remember reading a great technical analysis of Johann Bach's compositions, but you forgot to bookmark it (or would it be a *bachmark*?). You enter the keyword *bach* into the search bar and, because Johann Bach was a very important and famous composer, you get lots of results: too many, in fact. At first you're discouraged, but you remember that there's a site member who produces most of the site's good technical content, who's named *back2bach*. You see that his name is listed in the User facet, and there aren't many results in the facet count (the number in parentheses next to the facet). You click into the facet and quickly find the content you wanted.

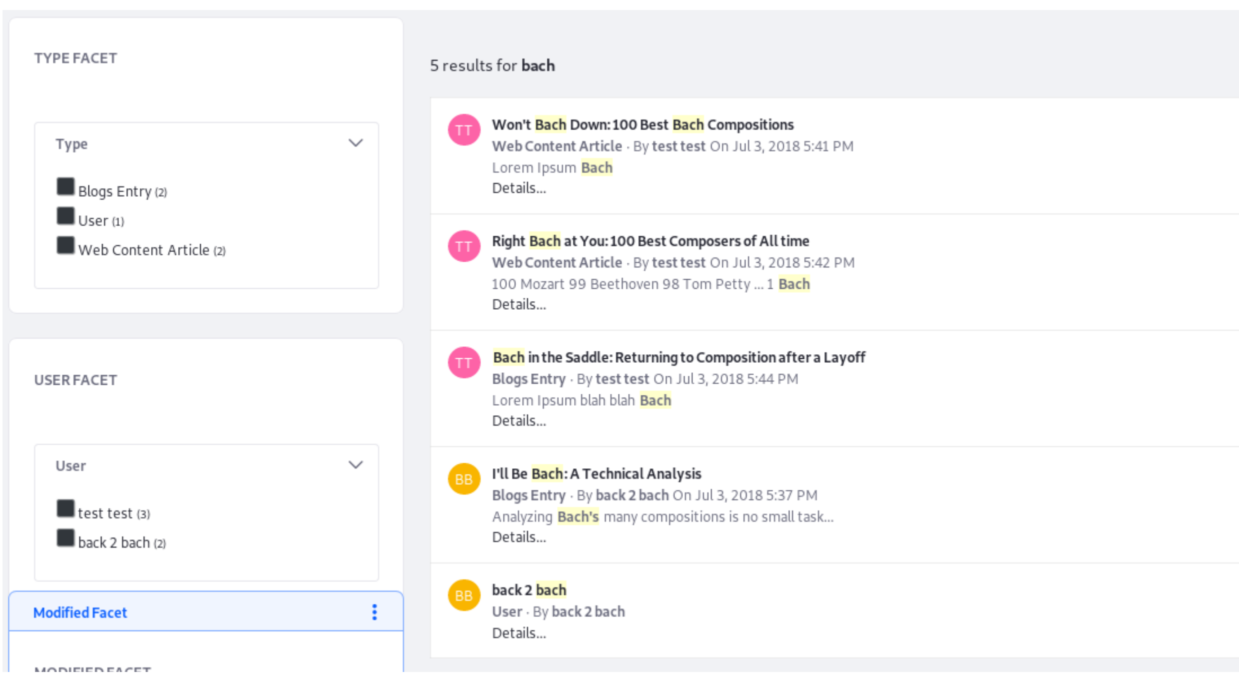

Figure 57.2: When presented lots of search results, facets narrow down the results list so users can find relevant content.

Clicking on a facet narrows down the search results. It's added to the filter list in the search query, and the results list is refined by the selected facets.

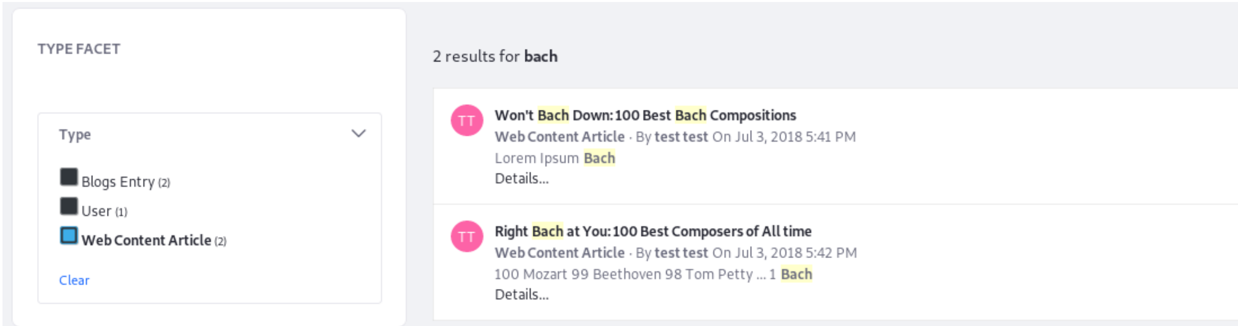

57.2 Multiple Facet Selection

Facet term selections within one facet are additive. Clicking more terms in the same facet expands the search results, because it's processed as if you want to see results matching *Term-1* OR *Term-2*, OR etc. To remove all the term selections from a facet, click the *Clear* link.

Figure 57.3: Facet terms are additive when applied in the same facet. Any Blogs Entry OR Web Content article matching the keyword is shown here.

Facet term selections from different facets are exclusive. Clicking facet terms from multiple facets narrows the results because they're processed as if you want to see results matching *Facet-1* AND *Facet-2*, AND etc. This is intuitive. The facets

Considering a case where you make two term selections in the Type Facet: (Blogs Entry and Web Content Article), and two term selections in the User Facet (James Jeffries and Marvin Smart). What results are displayed?

Blogs Entries OR Web Content Articles AND authored by James Jeffries OR Marvin Smart.

If Marvin and James each created four pieces of content (two blogs and two Web Content Articles), all eight would appear in the Search Results. Any Blogs or Web Content created by other Users are not shown, and assets of other Type created by Marvin and James are not displayed. Content that isn't Blog Entries or Web Content Articles created by other Users are obviously not searched.

Note: The new Search Facet widgets support the multiple selection of facet terms. Multiple facet selection is not supported in the classic Search portlet.

57.3 Facets and Friendly URLs

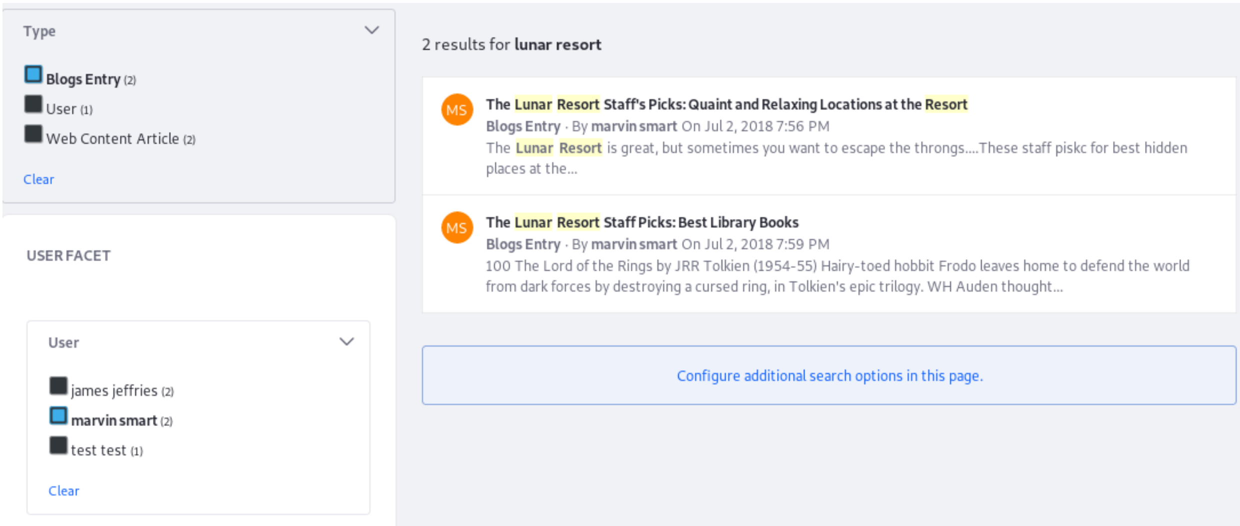

Figure 57.4: Facet terms selected from different facets are exclusive. These results must be of type Blogs Entry AND be from the User Marvin Smart.

In the classic, monolithic Search portlet, URLs like this were not uncommon:

```
http://localhost:8080/web/guest/home?_com_liferay_portal_search_web_portlet_SearchPortlet_formDate=1529671834606&p_p_id=com_liferay_portal_search_wel
site
```

The new Search functionality uses friendly search URLs for facet filtering. With the default settings, here's the default main search URL when searching for keyword *test*:

```
http://localhost:8080/web/guest/search?q=test
```

Selecting a facet term causes a new parameter to the above URL. For example, selecting *Blogs Entry* from the Type facet results in this URL:

```
http://localhost:8080/web/guest/search?q=test&type=com.liferay.blogs.model.BlogsEntry
```

Selecting another facet term from the same facet category appends the same parameter again, but with the newly selected value:

```
http://localhost:8080/web/guest/search?q=test&type=com.liferay.blogs.model.BlogsEntry&type=com.liferay.portal.kernel.model.User
```

The rest of the facets work the same way. Filtering by the last hour option in the Last Modified facet portlet produces this URL:

```
http://localhost:8080/web/guest/search?q=test&modified=past-hour
```

The parameter names are configurable for each facet.

Now that you know how facets work, read about configuring each of the included facets.

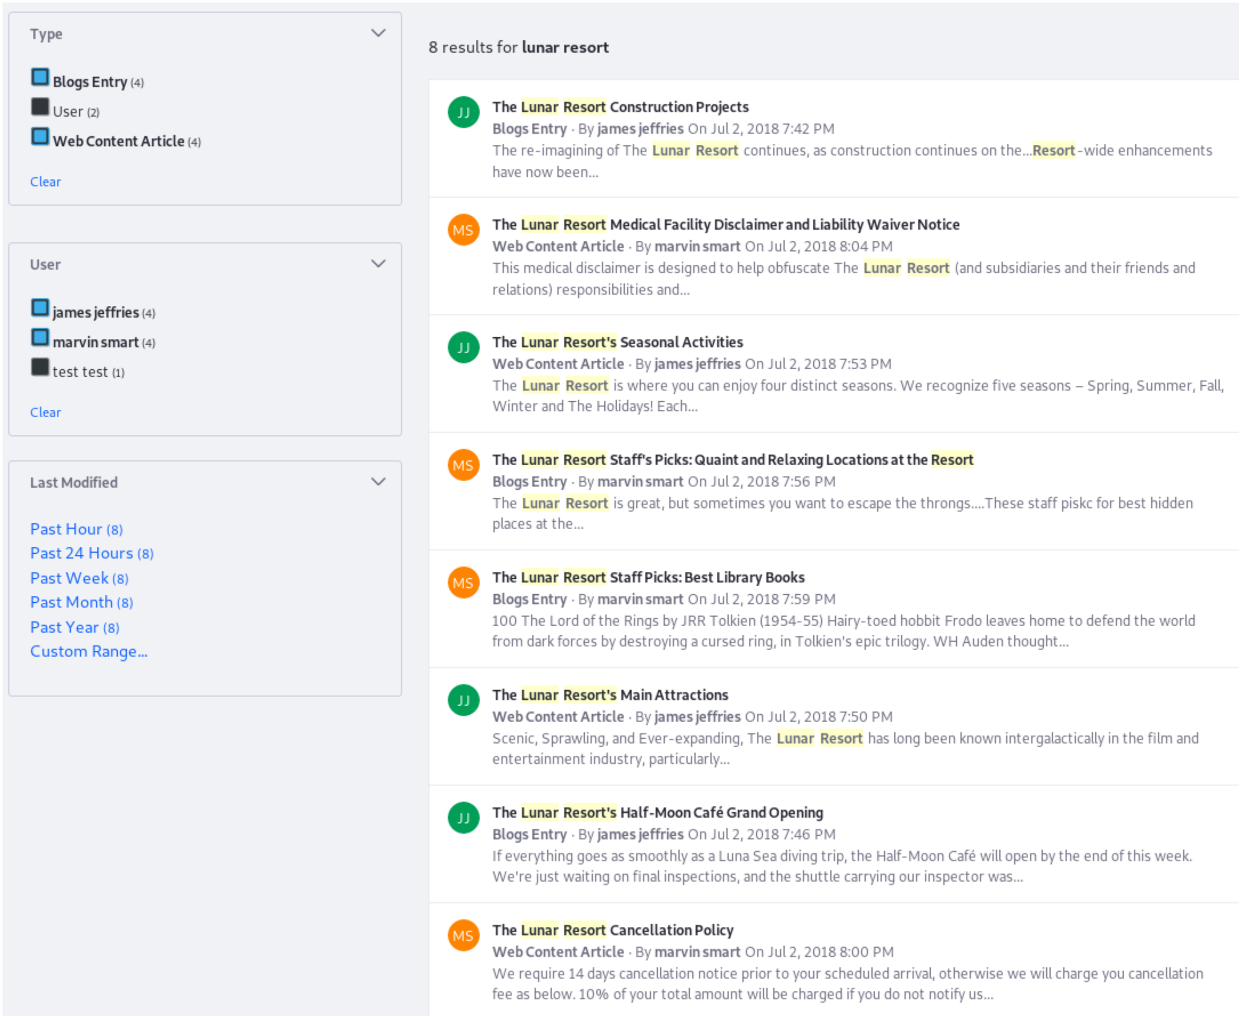

Figure 57.5: Both intra-facet and inter-facet selection is possible.

57.4 Site Facet

The Site Facet narrows search results down to those existing in a certain site. Each Site with content matching the searched keyword appears as a facet term.

For the Site Facet to display multiple sites, the Search Bar must be configured to search *Everything*. See more about search scope here. If not searching for Everything, only the current site is searched, and the Site Facet has nothing to display. When this occurs, the Site Facet is hidden on the page.

Note: Configuring the globally embedded page-top Search Bar to search for Everything not only configures the embedded Search Bar on all pages. It also ensures that the Search Page's Search Bar searches Everything, because the page-top Search Bar's configuration overrides the Search Page's Search Bar configuration. The same does not apply to other Search Bar widgets in the site. Each of these must be configured as desired.

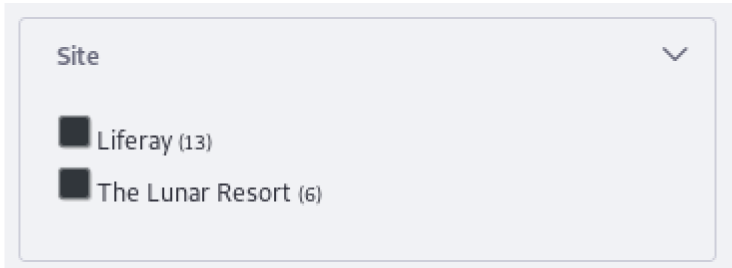

Figure 57.6: Each Site with matching content is a facet term.

If the global Search Bar is disabled, configure the Search Page's Search Bar widget to search for Everything.

To configure the search scope,

1. Open the Search Bar's Options menu () and click *Configuration*.
2. Set the Scope option to *Everything*.
3. Click *Save* and close the pop-up.

The Site Facet contains several configuration options:

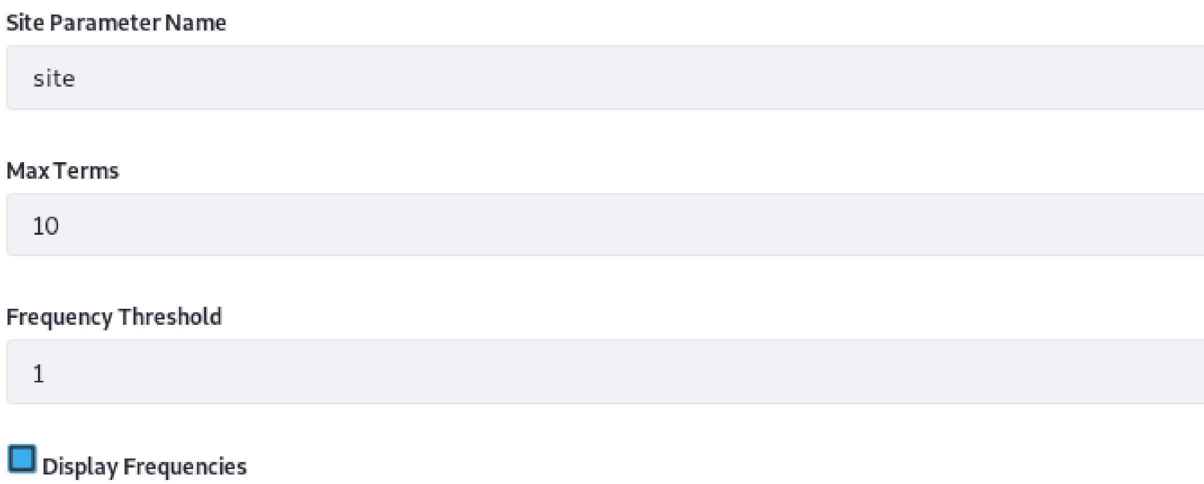

Figure 57.7: The Site Facet is configurable.

Site Parameter Name Set the URL parameter name for the Facet. The default is *site*. Searching for *lunar resort* and clicking on a site facet produces the URL

```
http://localhost:8080/web/guest/search?q=lunar resort&site=20126
```

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the list of facet terms. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether or not to display the term frequencies.

57.5 Type Facet

The Type Facet narrows search results down to those associated with a certain Asset Type. Each Type with content matching the searched keyword appears as a facet term.

Figure 57.8: Each Asset Type with matching content is a Type Facet term.

By default, all out of the box Asset Types are included as facet terms:

- Wiki Page
- Document
- User
- Blogs Entry
- Form Record
- Documents Folder
- Dynamic Data Lists Record
- Web Content Article
- Message Boards Message
- Calendar Event
- Knowledge Base Article

The Type Facet contains several configuration options:

Type Parameter Name Set the URL parameter name for the Facet. The default is *type*. Searching for *lunar resort* and clicking on a site facet produces the URL

```
http://localhost:8080/web/guest/search?q=lunar resort&type=com.liferay.blogs.model.BlogsEntry
```

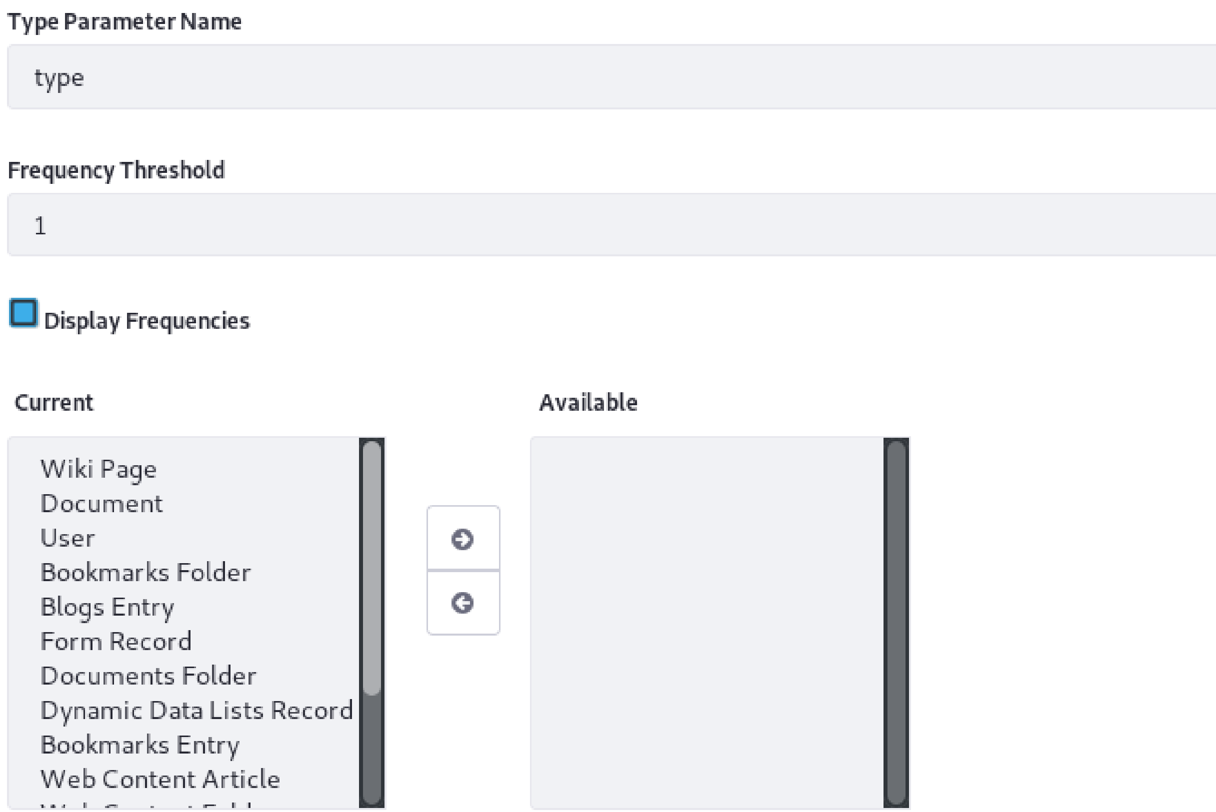

Figure 57.9: The Type Facet is configurable.

Frequency Threshold Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Display Frequencies Choose whether or not to display the term frequencies.

Current and Available Add or remove Asset Types from the facet. To remove types, select from the Current section by clicking and highlighting. Click the right arrow and move the Asset Type from *Current* to *Available*. Add Asset Types by moving them to the Current section.

57.6 Tag and Category Facets

If tags or categories were applied to an asset appearing in the result set, they're displayed in the Tag or Category facet, respectively. Like other facets with the Frequency Threshold configuration option, not all tags necessarily appear. By default the top 10 tags or categories are listed.

Tag and Category Facets contain identical configuration options:

Tag/Category Parameter Name Set the URL parameter name for the Facet. The default is *tag/category*. Searching for *lunar resort* and clicking on a *moon* Tag Facet term produces the URL

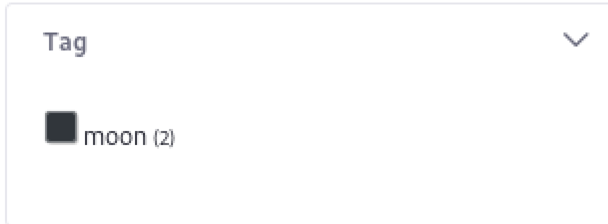

Figure 57.10: Each Tag or Category with matching content is a facet term.

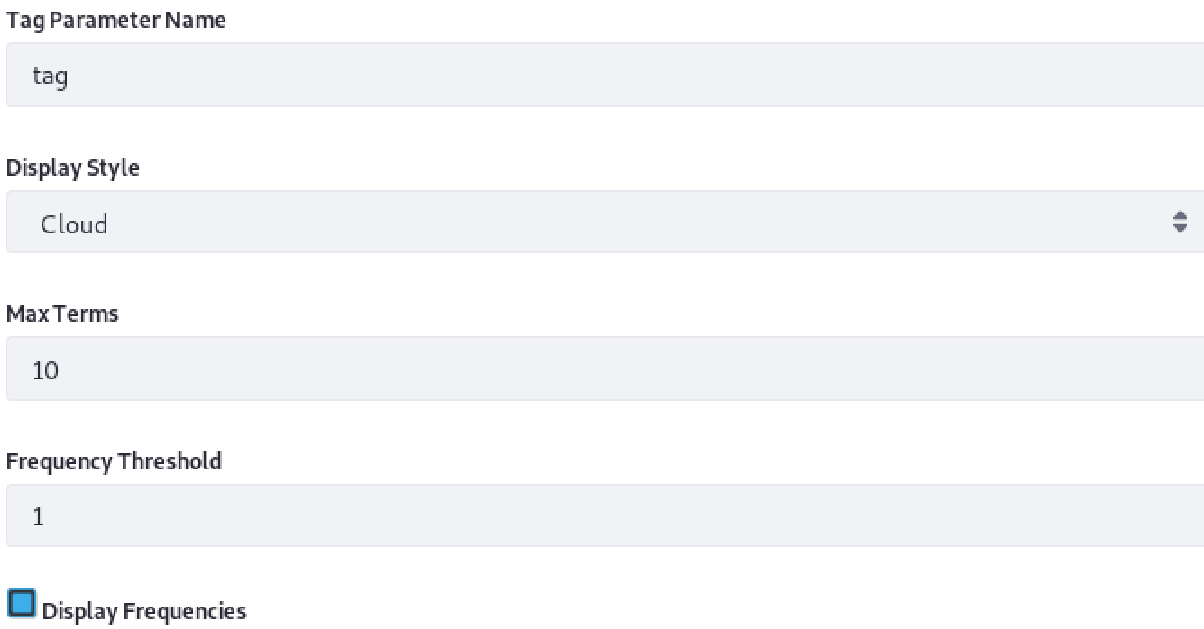

Figure 57.11: Tag and Category Facets are configurable.

```
http://localhost:8080/web/guest/search?q=lunar resort&tag=moon
```

Display Style Choose whether to display the facet terms in Cloud or List style.

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether or not to display the term frequencies.

57.7 Folder Facet

The Folder Facet narrows search results down to those contained in a certain Asset Folder. If you search for *space*, a Folder titled *Space Images* doesn't necessarily show up here. The content inside the folder must match the keyword. Only if its content matches the searched keyword does the Folder appear in the Folder Facet.

Folders of these Types appear as Folder Facet terms:

- Bookmarks Folder
- Documents and Media Folder
- Web Content Folder

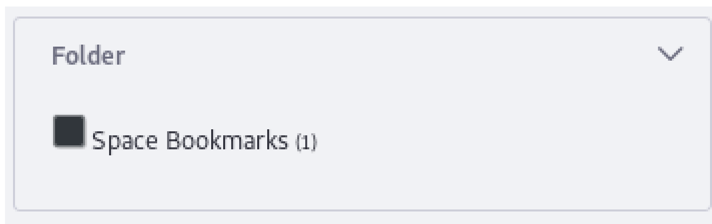

Figure 57.12: Each Folder with matching content is a facet term.

The Folder Facet contains several configuration options:

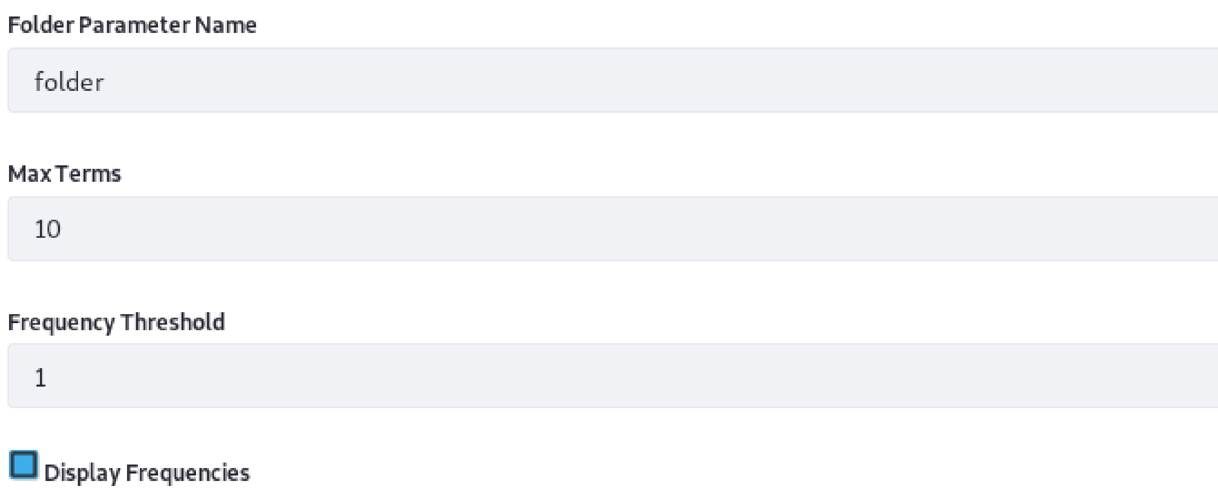

Figure 57.13: The Folder Facet is configurable.

Folder Parameter Name Set the URL parameter name for the Facet. The default is *folder*. Searching for *lunar resort* and clicking on a Folder Facet produces the URL

```
http://localhost:8080/web/guest/search?q=lunar resort&folder=38716
```

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether or not to display the term frequencies.

57.8 User Facet

The User Facet narrows search results down to those created by a certain User.

Figure 57.14: Each User with matching content is a facet term.

The User Facet contains several configuration options:

Figure 57.15: The User Facet is configurable.

User Parameter Name Set the URL parameter name for the Facet. The default is *user*. Searching for *lunar resort* and clicking on a User Facet produces the URL

```
http://localhost:8080/web/guest/search?q=lunar resort&user=38716
```

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the facet.

Display Frequencies Choose whether or not to display the term frequencies.

57.9 Modified Facet

The Modified Facet narrows search results down to those that match the searched keyword and that were created or modified during a certain time period.

Figure 57.16: Each time period with matching content is a facet term.

In addition to selecting a pre-configured time period, Users can select a Custom Range, specifying a From and To date using a date picker:

The Modified Facet supports configuration actions:

- Modify existing time ranges
- Delete existing time ranges
- Create new time ranges

Edit or create time ranges using a time range alias.

The available time range aliases include:

```
past-hour
past-24-hours
past-week
past-month
past-year
```

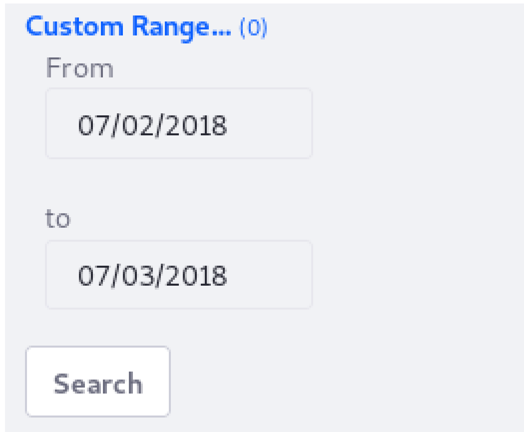

Figure 57.17: Users can include a Custom Range in the Modified Facet.

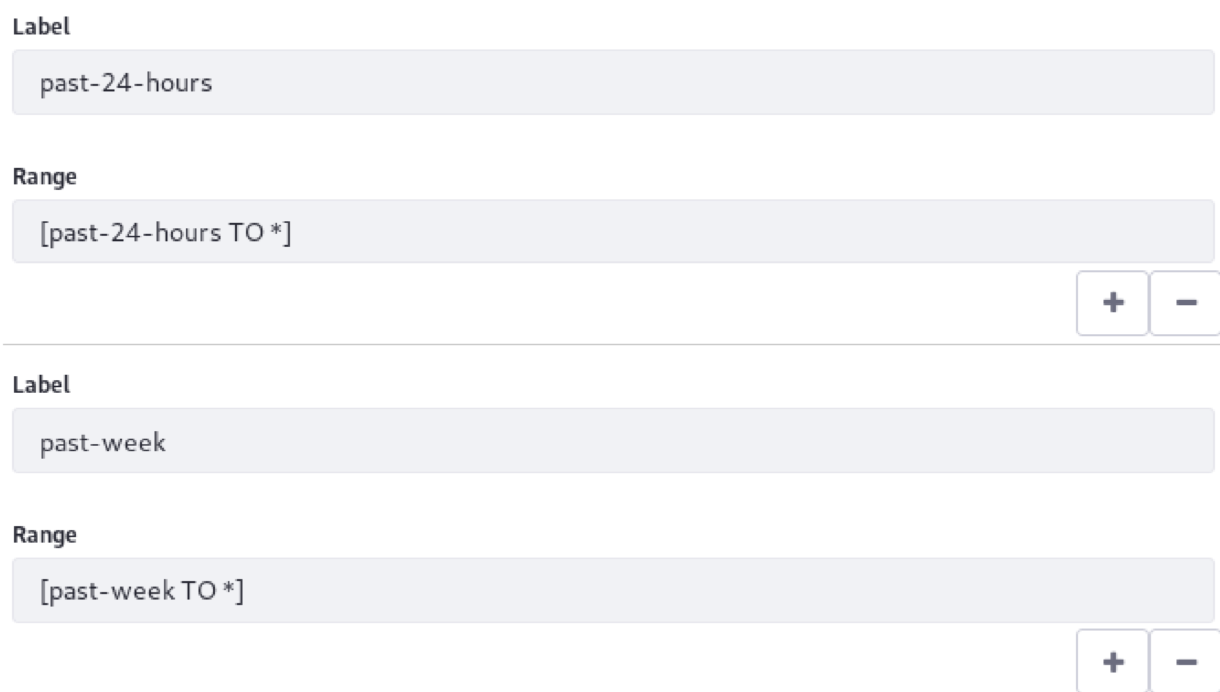

Figure 57.18: The time ranges are set in the facet's configuration.

Each Range has an alias and a Label.

By default, all the default ranges end in *, which evaluates to *now*. For example, the past-week range is

```
[past-week TO *]
```

You're not limited to ending Ranges. Instead of the *, specify another time range alias as the ending point.

To set up a range from 12 months ago to one month ago,

1. Click the plus button in one of the existing ranges.
2. Give it the label **1-12 Months Ago**.

Give it a Range value of

```
[past-year to past-month]
```

This gives you lots of flexibility in using alternative time ranges as Modified Facet terms.

57.10 Custom Facet

All Facets are configurable, allowing you to narrow down search results based on a shared characteristic (all Blog Entries, for example). The Custom Facet lets you create entirely new Facets. The first thing to do is enter the Custom Facet's configuration screen.

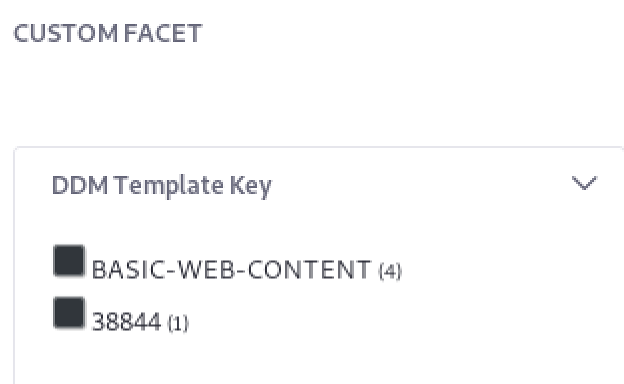

Figure 57.19: Custom Facets must be configured first.

The screenshot above shows a Custom Facet with the Web Content Article DDM Template Key as the field used to aggregate results. The next screenshot shows how it was configured.

Aggregation Field Specify the non-analyzed keyword field whose value is used to create the facet terms. If the value of the search result's `ddmTemplateKey` field is BASIC-WEB-CONTENT, that's what appears in the Custom Facet as a term.

Custom Heading Enter a human readable heading for the Custom Facet.

Figure 57.20: Configure a Custom Facet in no time.

Custom Parameter Name Set the URL parameter to use for the facet. With the configuration pictured above, searching for *lunar resort* and clicking on BASIC-WEB-CONTENT produces the URL

```
http://localhost:8080/web/guest/search?q=lunar&ddmtemplatekey=BASIC-WEB-CONTENT
```

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether to display the term frequencies.

Finding Indexed Fields

To use the Custom Facet, you must know which non-analyzed keyword field to specify. In most cases, if you know you require the Custom Facet, you know what information you need to create the facet:

1. Enable the Display Results in Document Form setting in the Search Results widget configuration.

2. Search for an asset that has the information you need.
3. Inspect the Document view of a resultant hit.

   Use the Key in the Custom Facet's *Aggregation Field*.

To browse the entire list of available fields, inspect the field mappings using your search engine's API.

For Elasticsearch, access the field mappings from your terminal using CURL to call the Get Mapping API:

```
curl -X GET "localhost:9200/_mapping/LiferayDocumentType"?pretty
```

Solr uses the ListFields API:

```
curl http://localhost:8983/solr/liferay/schema/
```

Here's a snippet of output from the Elasticsearch example:

```
"ddmStructureKey": {
    "store": true,
    "type": "keyword"
},
"ddmTemplateKey": {
    "store": true,
    "type": "keyword"
},
"defaultLanguageId": {
    "store": true,
    "type": "keyword"
},
"description": {
    "store": true,
    "term_vector": "with_positions_offsets",
    "type": "text"
},
"discussion": {
    "store": true,
    "type": "keyword"
},
```

Use Custom Fields to aggregate facet terms by shared non-analyzed keyword field values.

CHAPTER 58

SEARCH RESULTS

The ideal search experience involves a User entering a search term, waiting an infinitesimal amount of time, and having the perfectly matching asset delivered at the top of a list of other extremely relevant hits. Like this:

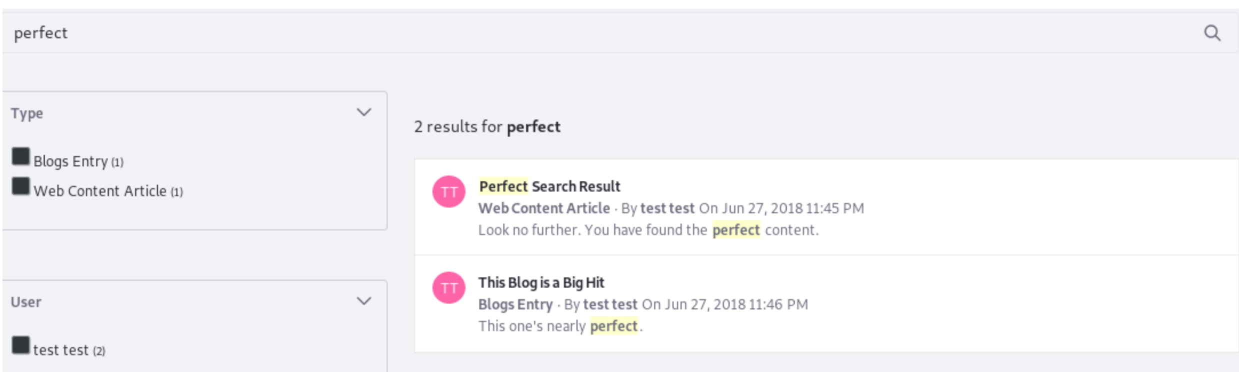

Figure 58.1: The goal is to return the perfect results to Users searching your site.

The developers of an asset control much about how the asset's information is stored in the search engine (this process is called *indexing*), and how its information is searched and returned in the search results. Developers who dislike how a particular asset behaves in search can use an *Indexer Post Processor* to modify the asset's indexing behavior and how search queries are constructed to look up the assets in Liferay DXP.

The Search Results behavior configurable through the UI is covered in this section:

- Search Results Display Settings
- Search Results Behavior

58.1 Display Settings

The Search Results widget's default display is a paginated list. Each list item is a summarized hit to a search query. Click on a specific result to look at it in more detail. Configure display options by opening the Search Results options menu () and selecting *Configuration*.

Enable Highlighting Highlight the search terms where they appear in the search result's title or summary.

Display Results in Document Form Display results as search documents. Never use this in production. Developers use this feature to view search responses in their indexed, document-based format. Part of a developer's job when writing search indexers is to convert documents (the objects that get indexed) to the actual object and back again. Thus, developers can see how their objects are being indexed. Once enabled, click the *Details...* link below the result summary to expand the result's document view.

Figure 58.2: Viewing a results document lets you inspect exactly what's being indexed for a particular asset. This is just a small portion of one document.

Display Selected Result in Context When an asset is clicked, show it in its native application. For example, if you click on a blog post in the search results, you see where the Blogs Entry is posted in the Blogs application. Note that you're not in the search context after clicking on a search result. When this option is unchecked, the asset displays in an Asset Publisher window while still in the search context. If you have the right permissions, you can even edit the content directly from the Search context. Click the back arrow to return to the search results.

The next three configurations concern results pagination.

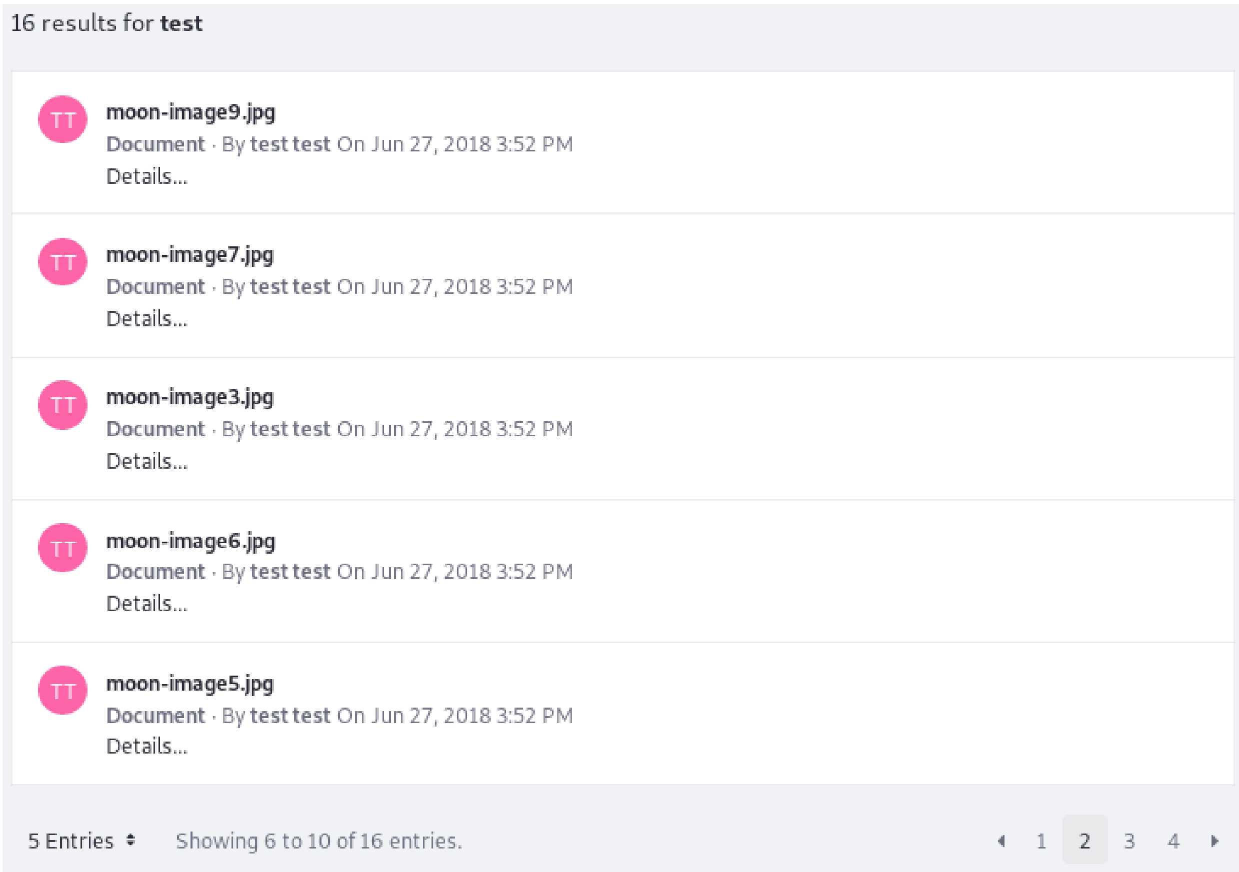

Figure 58.3: The number of results per page and the URL parameter names used to control pagination behavior are configurable.

Pagination Start Parameter Name Set the name of the URL parameter for the results page. If the default value *start* is preserved, this URL displays when the User navigates to the second results page after searching for *test*:

```
http://localhost:8080/web/guest/search?q=test&start=2
```

Pagination Delta Set the number of results to display per results page. Defaults to *20* unless you customized the `search.container.page.default.delta` property in your `portal-ext.properties` file.

Pagination Delta Parameter Name Set the name of the URL parameter that stores the Pagination Delta value. This becomes visible in the browser if the User changes the number. If the User selects 10 results per page and searches for *test*, the Search Page is reloaded with this URL:

```
http://localhost:8080/web/guest/search?q=test&delta=10
```

For further reading, check out how to return suggestions for better search terms (for example, "Did you mean...") when not enough results are returned initially.

58.2 Search Results Behavior

The previous article covered ways to display search results. This article covers these additional Search Results concepts and configurations:

- Filtering search results with facets
- Understanding search results relevance
- The effect of permissions on search results
- Search results in the staging environment
- Search results summaries
- Search results term highlighting

Filtering Results with Facets

Results are filtered using *facets*. Most users have encountered similar filtering capabilities in other applications, particularly during commerce activities. Users enter a search term, are presented with a list of results and search facets, which you can think of as buckets that group results together if they share a common characteristic.

Administrators can configure facets. Read about configuring facets to learn more.

Search Results Relevance

The search engine decides which results appear at the top of the list using the concept of *relevance*. Relevance is a score calculated by the search engine. There are numerous factors contributing to the total score of a returned document, and all of the implementation details of how relevance scoring works are algorithms provided by the search engine.

Permissions and Search Results

Users lacking VIEW permission on an asset don't see it in the search results. A logged in User with the Site Administrator role likely sees more search results than a guest User to the site.

In the background, there are two rounds of permissions checks. The first permissions check, *pre-filtering*, happens in the search engine's index. It's faster than checking database permissions information, but occasionally the search index can have stale permissions information. To ensure the search engine's index has correct, up-to-date permissions information, a second, last-second permissions check, *post-filtering*, is performed on the results prior to their display.

Initial Permissions Checking

The first round of search results permissions filtering adds filter clauses to the search query. This ensures that results return from the search engine pre-filtered, containing results the current User can view.

This initial permission checking is configurable at *Control Panel* → *Configuration* → *System Settings* → *Search* → *Permission Checker*. It includes two system level settings to configure how search processes User permissions.

Include Inherited Permissions Ignore this setting. It's deprecated, no longer used anywhere, and will be removed in a future release.

Permissions Term Limit Limits the number of permission search clauses added to the search query before this level of permission checking is aborted. Permission checking then relies solely on the final permission filtering described below.

The only reason to limit permissions terms is performance. Users with administrative access to lots of sites and organizations generate many permissions terms added to the query. Too many terms in a query can make the search engine time out.

Final Permissions Checking

A final round of permission checking happens prior to presenting results in the UI. For example, the User searches for *liferay*, and the search engine returns all relevant forum posts. As the Search Results iterates through the list of relevant forum posts, it performs one last permission check of the post to ensure the User can view the post and its categories. If a matching forum post exists in a category the User doesn't have permission to view, it isn't displayed in the list of search results.

This final round of permission checking is configurable at *Control Panel* → *Configuration* → *System Settings* → *Search* → *Default Search Result Permission Filter*. It includes two settings:

1. The first setting, Permission Filtered Search Result Accurate Count Threshold, specifies the maximum number of search results to permissions-filter before results are counted. A higher threshold increases count accuracy, but decreases performance. Since results in the currently displayed page are always checked, any value below the search results pagination delta effectively disables this behavior.

2. The second setting, Search Query Result Window Limit, sets the maximum batch size for each permission checking request. . This is again impacted by pagination. For example, if there are 100 results per page, and a User wants to jump all the way to page 200 of the search results, all results between page one and 200 must be checked to ensure the User has permission. That's 20,000 results to permissions check. Doing this in one trip to and form the search engine can result in performance issues. Set the maximum batch size for each permission checking request.

Search and Staging

With staging, content is placed first in a preview and testing environment before being published for consumption by end Users (on the live site). Content added to the search index is marked so that the search API can decipher whether an item is live or not. In the live version of the site, only content that's marked for the live site is searchable.

In the staged version of the site, all content—live or staged—is searchable.

Result Summaries

A result summary includes the information from a document that the asset's developer felt is most useful to end Users searching for the asset. Each asset can have different fields included in the summary. For assets with text content, a common summary format includes the *title* and some of the *content*, with title displayed first. The asset type always appears on the second line, and a snippet of the content that matches the search term is on the last line. Assets without content fields, like Documents and Media documents, display the description instead.

Users are different. Only the User's full name and the asset type (User) appear in User result summaries.

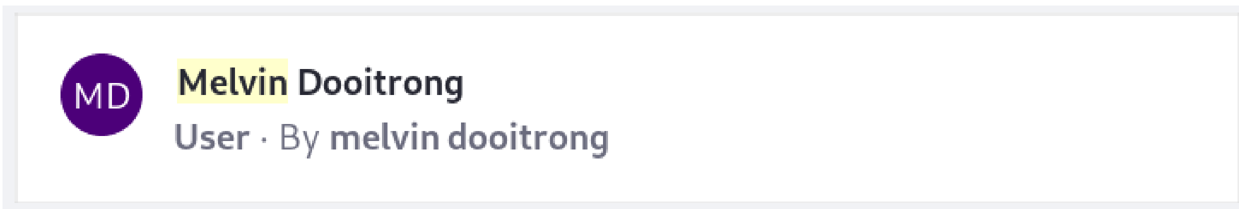

Figure 58.4: User summaries contain only the User's full name.

For assets that contain other assets (Web Content, Documents & Media, and Bookmarks folders) or whose content is not amenable to display (Dynamic Data List Records and Calendar Events), it makes more sense to display the title, asset type, and description in results summaries. There'd never be anything in a content field for these assets.

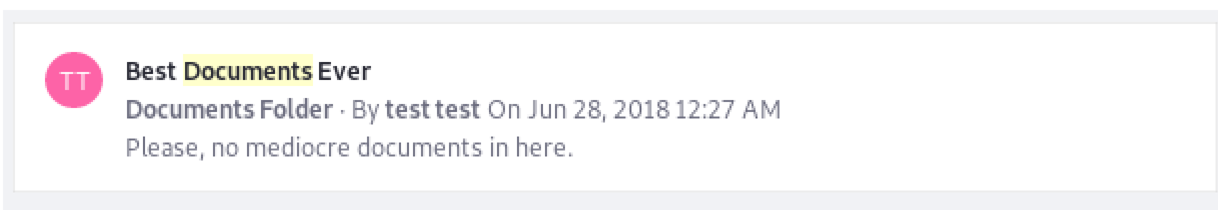

Figure 58.5: Documents and Media, Web Content, and Bookmarks folders include titles and descriptions in their summaries.

Bookmarks entries show the title and the URL.

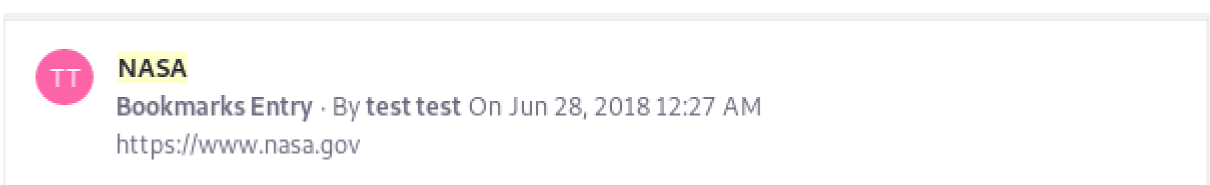

Figure 58.6: Bookmarks Entries summaries show the title and the URL.

The asset developer determines which fields are summary-enabled, but there's logic invoked at search time that determines precisely the part of the summary fields to display. For example, a `content` field can have a lot of text, but the summary doesn't show it all. Instead, it shows a relevant snippet of the field's text. If the keyword searched for is present in the summary field, that portion of the field is used in the summary. In addition, the matching keyword is highlighted in the summary.

Highlighting

By now you've probably noticed that search terms appearing in the summary are highlighted by default. If this is undesirable, disable it in the widget configuration screen.

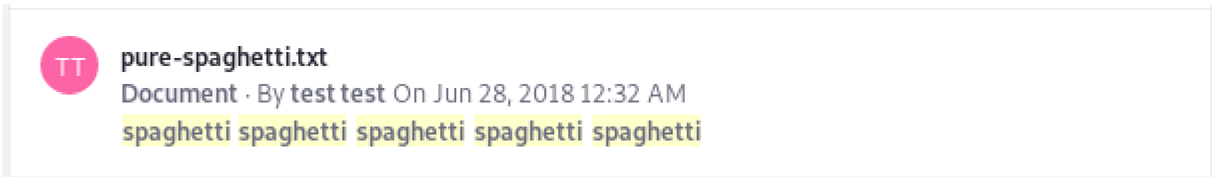

Figure 58.7: Some document summaries have lots of highlights if the search term matches text that appears in the summary.

Highlighting is a helpful visual cue that hints at why the result is returned, but beware. A hit can score well and thus be returned near the top of the results, without having any highlights in the summary. That's because not all indexed fields appear in the summary. Consider a User named Arthur C. Clarke. He has an email address of *acc@authors.org*, which is searchable. Because results summaries for Users only contain the full name of the User, searching for Mr. Clarke by his email address returns the User, but no term is highlighted.

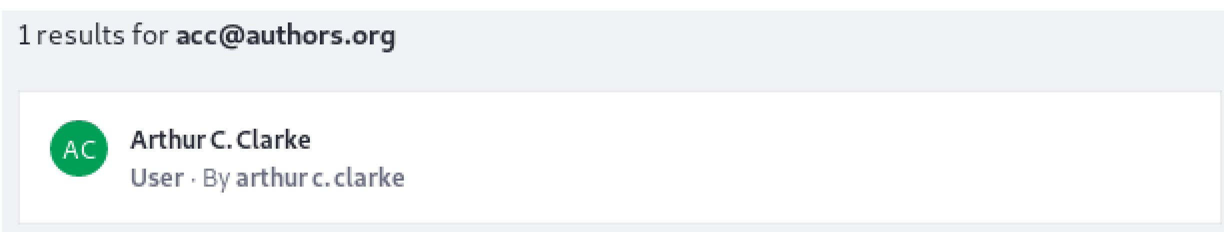

Figure 58.8: Results that match the search term won't always have highlights.

There are additional cases where search results don't have highlighting.

58.3 Configuring Search

Configuring Search could mean lots of different things:

- System scoped search configuration
- Reindexing to make sure the search indexes are current with the database
- Tweaking the search widgets added to pages
- Creating new Search Pages
- Configuring the connectors that let Liferay DXP and the search engine communicate

In fact, *Configuring Search* means all those things. This is a high level overview of what search behavior is configurable out of the box, and importantly, *where* to find search configuration options.

System Scoped Search Configuration

System scoped search configurations are primarily found in System Settings.

1. Go to *Control Panel* → *Configuration* → *System Settings*.
2. Click the *Search* category under the Platform section.
   Alternatively, search for *Search*.

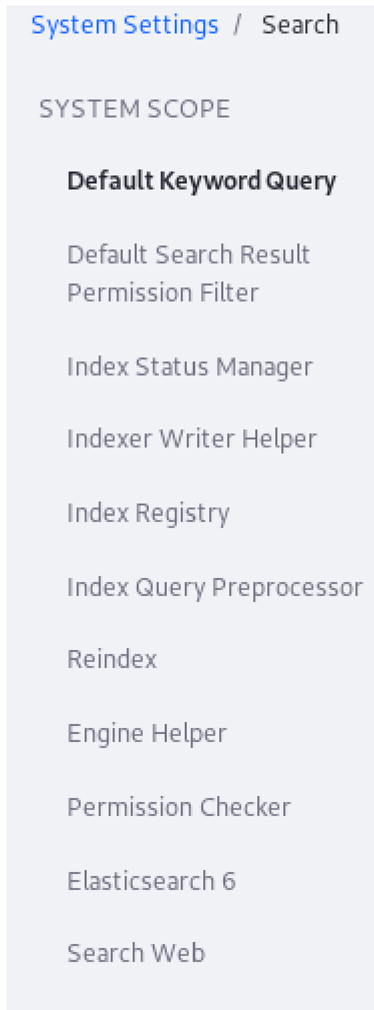

Figure 58.9: There are numerous system scoped entries for search in System Settings.

These system scoped configurations are available in System Settings:

Default Keyword Query

The Default Keyword Query entry contains one setting:

`disabledEntryClassNames`: The `DefaultKeywordQueryContributor` code automatically adds `description`, `userName`, and `title` fields to the keyword search query. Specify the entry class names `DefaultKeywordQueryContributor` should ignore.

Configuration file: `com.liferay.portal.search.configuration.DefaultKeywordQueryConfiguration.config`

Default Search Result Permission Filter

The Default Search Result Permission Filter entry allows configuration of *post-filtering permission checking* (database permission checking that occurs after the results are returned from the search index). Read here for more information on these settings:

- `permissionFilteredSearchResultAccurateCountThreshold`
- `searchQueryResultWindowLimit`

Configuration file: `com.liferay.portal.search.configuration.DefaultSearchResultPermissionFilterConfiguratio`

Index Status Manager

The Index Status Manager entry has one setting:

`indexReadOnly`: Suspends all indexing operations and writes to the search engine. Searches return only the documents already indexed. This is useful for speeding up large data imports, but it should be disabled and a full re-index executed once the import is finished.

Configuration file: `com.liferay.portal.search.configuration.IndexStatusManagerConfiguration.config`

Indexer Writer Helper

The Index Writer Helper entry contains only one valid entry. The second, `indexReadOnly`, is deprecated and unused, so setting it has no effect. Use `indexReadOnly` from the Index Status Manager instead.

`indexCommitImmediately`: When *true* (the default), each write request forces the search engine to refresh the index reader, potentially flushing transactions to disk. This may negatively impact search engine performance. The default behavior is to commit immediately for index writing on individual assets (e.g. add blog, update blog) but delay commits for bulk index writing operations (e.g. index all users, index all form entries) until all entries have been sent to the search engine. Setting this to false changes the behavior for individual index operations, and may cause applications like Asset Publisher to exhibit a delayed response when showing newly added content. See the Elasticsearch documentation for more information.

Configuration file: `com.liferay.portal.search.configuration.IndexWriterHelperConfiguration.config`

Index Registry

Configure the buffering of index requests:

`buffered`: Disable or configure the buffering of indexing requests. To stop the buffering of index requests, choose *Disabled*.

`maximumBufferSize`: If buffering is enabled, set the Maximum Buffer Size so that any additional indexing requests are executed immediately.

`minimumBufferAvailabilityPercentage`: When the capacity of the buffer has only the specified percent of space left, the existing requests in the buffer are executed in one batch and removed from the buffer.

Configuration file: `com.liferay.portal.search.configuration.IndexerRegistryConfiguration.config`

Index Query Preprocessor

This entry has one repeatable property (use array syntax if you're defining via OSGi configuration file):

`fieldNamePatterns`: Fields with names matching the patterns set here are treated as non-analyzed keyword fields. Instead of scored full text queries, matching is performed by non-scored wildcard queries. This is a resource intensive operation that degrades search engine performance as indexes grow larger. For substring matching, relying on the ngram tokenizer usually performs better.

Configuration file: `com.liferay.portal.search.configuration.QueryPreProcessConfiguration.config`

Reindex

This entry contains only one property:

`indexingBatchSizes`: Set the number of documents indexed per batch for model types that support batch indexing—defaults to `com.liferay.journal.model.JournalArticle=10000`. For models with large documents, decreasing this value may improve stability when executing a full re-index.

Configuration file: `com.liferay.portal.search.configuration.ReindexConfiguration.config`

Engine Helper

This entry has one repeatable property (use array syntax if you're defining via OSGi configuration file):

`excludedEntryClassNames`: Exclude an asset type from being searched in the catchall query constructed for the Search application. For example, fields of the Organization asset must be indexed to be searchable from the Users and Organizations application, but should not be searched in the Search application. Thus, Organizations are added to `excludedEntryClassNames`.

Configuration file: `com.liferay.portal.search.configuration.SearchEngineHelperConfiguration.config`

Permission Checker

Configure *pre-filtering permission checking* (permission checking on the search index) behavior. See here for more information on these properties:

- `includeInheritedPermission`
- `permissionTermsLimit`

Configuration file: `com.liferay.portal.search.configuration.SearchPermissionCheckerConfiguration.config`

Elasticsearch 6

Configure the connection between Liferay DXP and Elasticsearch 6. See here for more information on these properties:

- `clusterName`
- `operationMode`
- `indexNamePrefix`
- `numberOfIndexReplicas`
- `numberOfIndexShards`
- `bootstrapMlocakAll`

- `logExceptionsOnly`
- `retryOnConflict`
- `zenDiscoveryUnicastHostsPort`
- `networkHost`
- `networkBindHost`
- `networkPublishHost`
- `transportTcpPort`
- `transportAddresses`
- `clientTransportSniff`
- `clientTransprtIgnoreClusterName`
- `clientTransportPingTimeout`
- `clientTransportNodesSamplerInterval`
- `HttpEnabled`
- `HttpCrsEnabled`
- `HttpCorsAllowOrigin`
- `HttpCorsConfigurations`
- `additionalConfigurations`
- `additionalIndexConfigurations`
- `overrideTypeMappings`
- `synchronizedSearch`

Configuration file: `com.liferay.portal.search.elasticsearch6.configuration.ElasticsearchConfiguration.confi`

Search Web

This entry contains one property: `classicSearchPortletInFrontPage`: Revert the default search experience from using the new Search Widgets to the classic Search Portlet that was standard in past releases. See here for more information.

Configuration file: `com.liferay.portal.search.web.internal.configuration.SearchWebConfiguration.config`

Reindexing from Search Administration

In addition to the System Settings for Search, the action of recreating the search indexes is a system scoped action.

1. Go to *Control Panel* → *Configuration* → *Search*.

2. Re-index one of these:

    - All indexable assets
    - An individual indexable asset
    - All spell check indexes

Portal properties are system scoped configurations as well. The Lucene Search portal properties configure low level search behavior. Review the properties and their descriptions and determine if they apply to your search requirements.

Site Scoped Search Configuration

Search isn't configurable at the Site Scope by the strict definition of Site Scoped Configuration. However, Search Pages influence site-specific search behavior. Commonly, Search Pages contain search widgets configured to search for all content within a particular Site.

In addition, the Header Search (the Search Bar embedded in every Site page by default), whether populated by the new Search Bar widget or the legacy Search portlet, is Site scoped. Only one instance of the Header Search application exists per Site, and configuring it in one page context configures it for the entire Site.

Because of the modularity of Search, there are some important configuration nuances to be aware of when using the new Search widgets:

- If the Header Search uses the Search Bar widget, its configuration always requires a *destination page* to be set, where Users are redirected to complete their search activity, interacting with the other Search widgets (Results, Facets, Suggestions etc.). Search destination pages are ordinary pages holding the Search widgets. You can have as many pages with Search widgets across the Site as you want.
- Unlike the legacy Search portlet, the new Search Bar widget is instanceable, so one page can contain multiple Search Bar widgets configured differently. All Search Bar instances must point to a Search Page to be effective.
- **Important**: if the destination Search Page has a Search Bar widget instance besides the embedded Header Search, the configurations of the Header Search take precedence over the page's widget instance.

  Conversely, searching from a Search Bar widget instance on other pages honors their configurations, even if they differ from the Header Search configuration.

Note: On Liferay Portal 7.1 CE GA1 and Liferay DXP 7.1, you must configure both the Search Page's Search Bar and the Header Search with matching configurations to achieve consistent search behavior (otherwise either one may take precedence.) This has been fixed in LPS-83193 and will be available on the next Fix Pack or GA release.

See the documentation on configuring of a Search Bar for more information.

Widget Scoped Search Configuration

Several search widgets are available, and each one has its own configuration options:

Search Results Configure how search results are displayed. Read here for more information.

Search Bar Configure the behavior of how search keywords are processed. See here for more information.

Search Facets Configure each facet's behavior and URL parameters. See here for more information.

Search Options This is a special case, where configuring this widget defines page scoped behavior. Add the Search Option widget to a page and define two booleans for the Search Page:

- Allow Empty Searches: By default, failure to enter a keyword returns no results. Enabling this ensure that *all* results are returned when no keyword is entered in the Search Bar.

- Basic Facet Selection: By default, facet counts are recalculated after each facet selection. Enable this to turn off facet recounting.

Search Suggestions Suggest better queries and spell check queries. See here for more information.

Search Insights Add this to the Search Page to inspect the full query string that's constructed by the back-end search code when the User enters a keyword. Only useful for testing and development.

Chapter 59

Forms

Users sometimes must give you information so you can help them. Whether you're asking them to submit a brief survey or apply for a mortgage, you must design a form. Liferay Forms gives you almost limitless form building capability. For a complete list of the form fields available, visit the form fields reference article.

Because the complexity of use cases for Forms varies from a single input field to many pages of fields with different configurations, it makes sense to show you how to build and publish simple forms very quickly, and then show you all the additional features you can use for more complex use cases. Here's a sampling of the what the Forms application can do:

- Populate a Select or Radio field with a REST Data Provider
- Make a field appear based on the value of another field
- Add extra pages to the form
- Enable CAPTCHA for a form
- Store results in JSON
- Enable workflow for the form
- Redirect to a different URL after a successful form submission
- Send an email notification to administrators whenever a form is submitted
- Provide a default value (entered if left alone by the user) or a placeholder value (not entered if left alone by user) for each field
- Validate fields using a number of different criteria
- Redirect users to a success page after form submission
- Define Form Rules to create dynamic form behavior (for example, show or hide a field based on input in another field).
- Translate form text into any supported language.
- Create partial forms (with fields and other elements and specific configurations) and save them for reuse.
- Drag and drop fields onto the form layout.
- Duplicate a form instead of starting a similar form from scratch.

Despite this long list of more complex options, developing a simple, elegant form to suit basic needs takes little effort. The next article covers basic form building.

59.1 Forms and Lists

When you need a form, what you're really looking for is data. There are two applications for building forms to collect precisely the data you need:

1. Liferay Forms: The primary form building application is for the simplest one or two question survey to the most complex, multi-page, homeowners insurance application containing rules and lists populated by a REST data provider.

2. Dynamic Data Lists (DDL): Provides a user interface tool for building reusable form- and list-based applications intended for display on pages, using templates.

Kaleo Forms: If you're a Liferay Digital Enterprise customer, there's a third form building tool called Kaleo Forms. It integrates form building with workflow to create form-based business processes, like a Conference Room Checkout Form, or a Support Ticket Process so support tickets go through the proper channels on their way to resolution. Read more about Kaleo Forms in the workflow section.

59.2 Which Form Builder Should I Use?

Liferay Forms (also referred to as Forms) is a relatively new application, first appearing in Liferay DXP version 7.0. If you can use Liferay Forms for your use case, you should.

So the question "Which form builder should I use?" can be restated to "When should I use Dynamic Data Lists?"

- Use Dynamic Data Lists (DDL) if you need a way for users to enter data, *and* you need to display the data in the user interface.

- Use DDL if you need to style your lists and forms with templates.

- Use DDL if there's a field type you need that's not included (yet) in Liferay Forms. These are the field types included in DDL that *are not* in Liferay Forms at the time of this writing: - Geolocation - Web Content - Link to Page

It's important to note that these (and more!) form field types will be included in future versions of the Liferay Forms application.

When all form building features are fully merged into Liferay Forms, the best features of DDL, all the new features of Liferay Forms, and all future improvements will be in one application. Now is the time to familiarize yourself with Liferay Forms and begin using it for all your form building needs, except for the narrow use cases described above.

59.3 What's New with Liferay Forms

Liferay Forms is evolving. This article compiles the prominent changes and additions to Forms. More details on these features are found in the other articles of this section.

Form Rules Set dynamic form behavior by specifying conditions (*if this happens in field X*) and actions (*then do this: X*). Available actions include Show, Enable, Require, Autofill, Calculate, and Jump to Page.

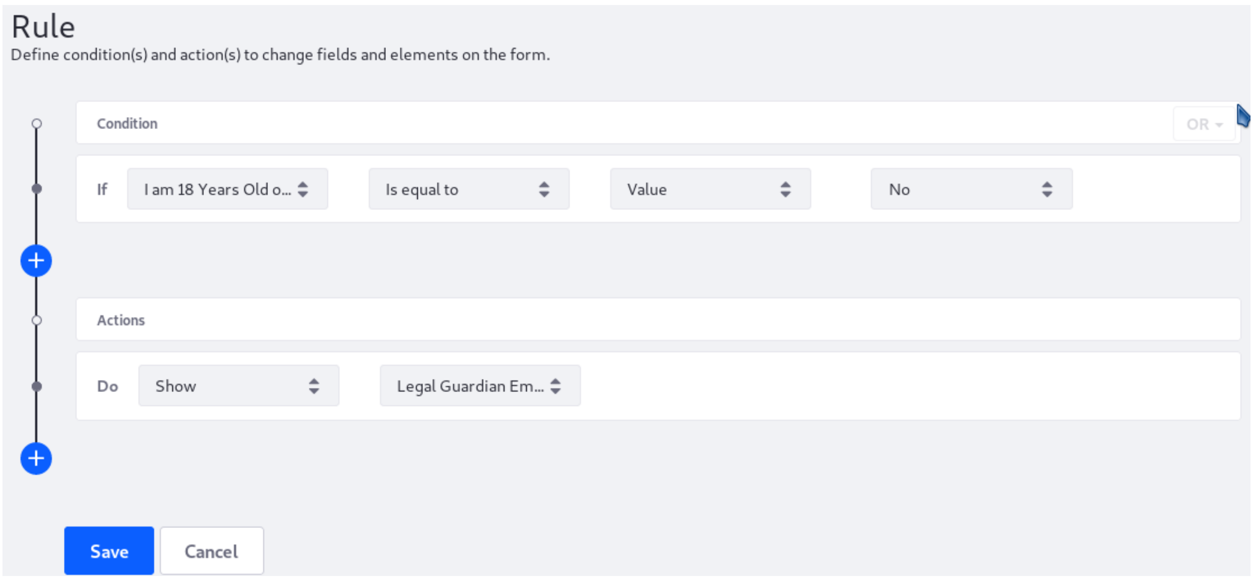

Figure 59.1: Form Rules trigger an action in response to some condition.

Form Translation Provide translations of the same form into the platform's supported languages. Form Users select the language they'd like to see.

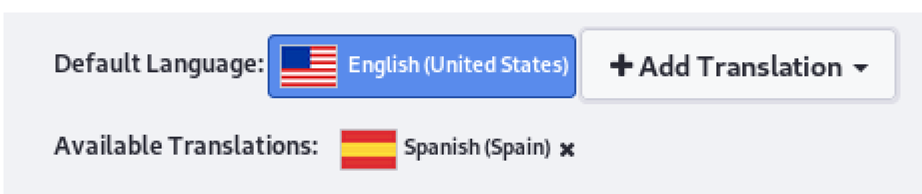

Figure 59.2: Translate a Form into a supported language.

New Fields and Properties File Upload fields let logged in Users upload documents to the Form. Numeric fields validate that input is numeric. Grid fields let Users select from options laid out in rows and columns. One selection can be made per row.

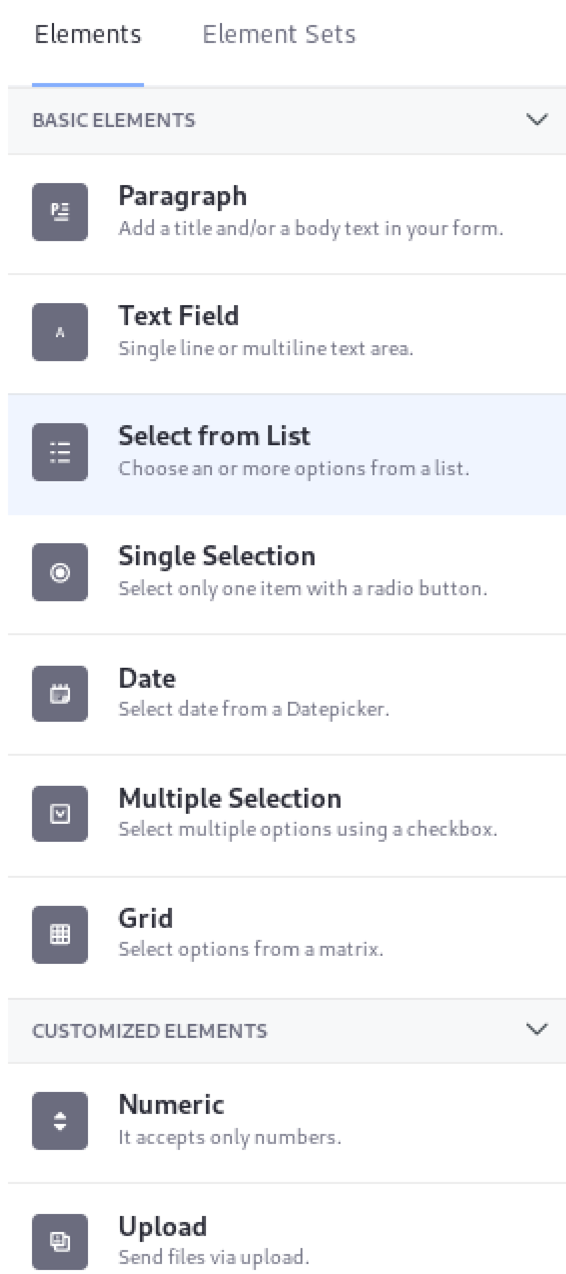

Figure 59.3: File Upload, Grid, and Numeric fields are new.

Element Sets Create reusable Form fragments that can be used to quickly populate new Forms with common fields and configurations.

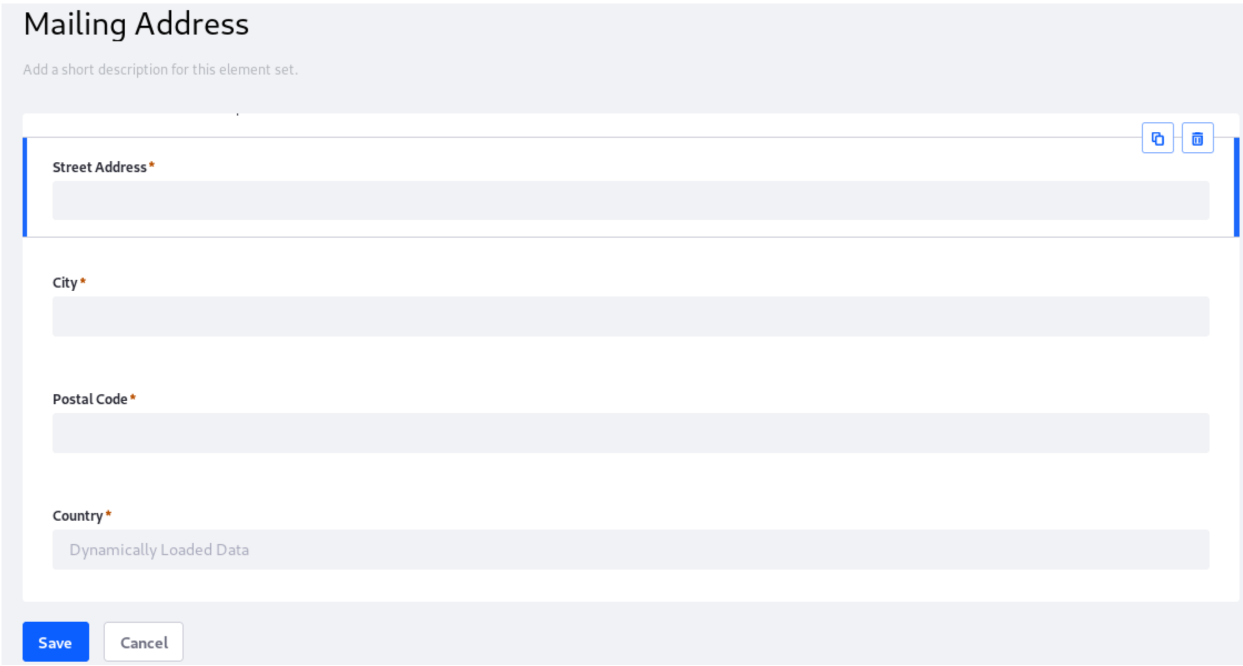

Figure 59.4: Reusable Element Sets are great for quickly populating new forms.

Autosave By default, a Form is auto-saved every minute (without the Form builder clicking *Save*). The duration is configurable in Control Panel → Configuration → System Settings → Forms. Auto-saves do not propagate to published Forms.

Autocomplete Text Fields Using a Data Provider or manually typed options, configure a Text field to offer Users autocompleted values to choose from as they type in the Text field. Set up autocomplete in the Text field's Properties.

Form Success Page Add a Form Page that displays after a successful Form submission.

Duplicate Forms and Form Fields Duplicate a Form to avoid building a similar one from scratch. Duplicate Form Fields if you need multiple similarly configured fields.

Validate Text Field Entries Add validation logic to a Text field and create custom messages for form validation errors. Use regular expressions in your validation by choosing the *Does not match* value in the *If Input* Select from List field.

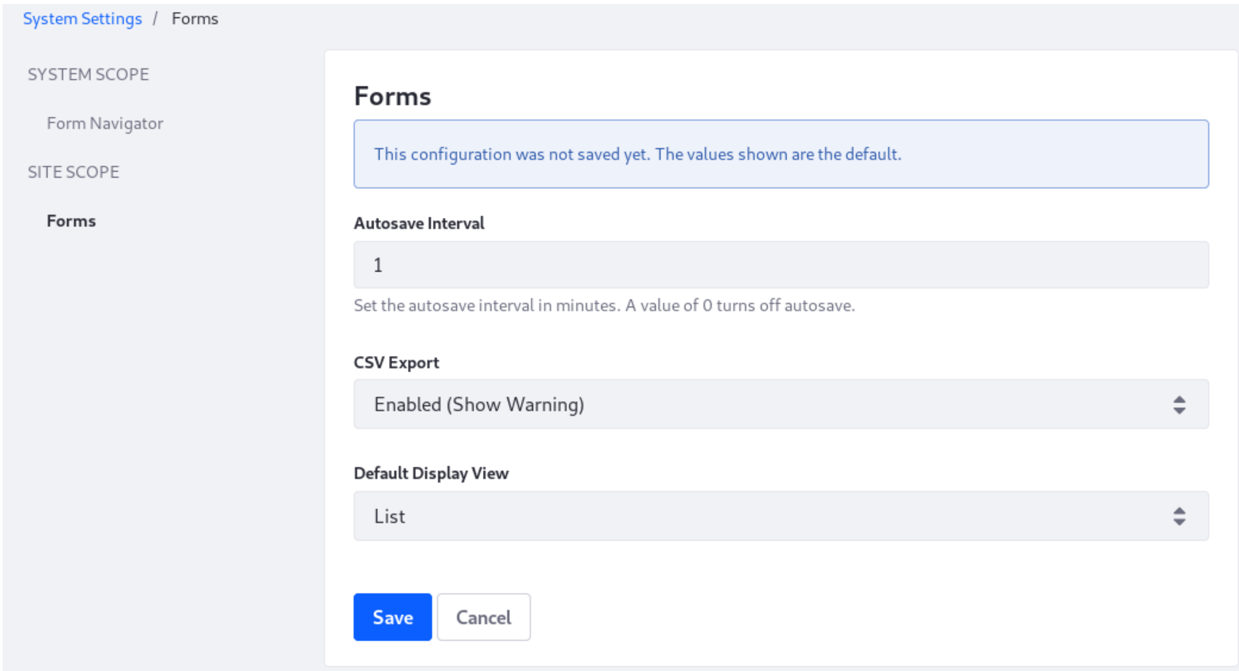

Figure 59.5: Autosaving Forms ensures your recent changes are preserved.

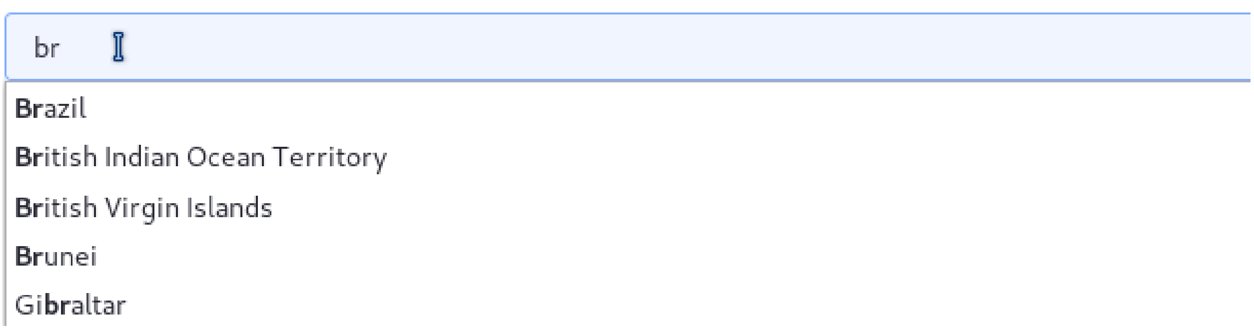

Figure 59.6: A User beings typing in a Text field and is pleasantly surprised by a list of autocomplete options.

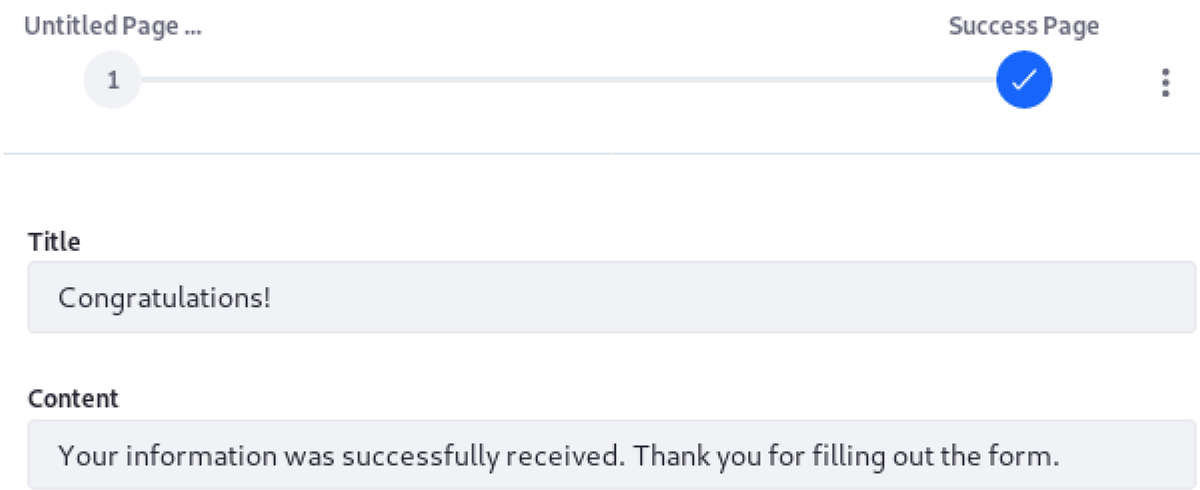

Figure 59.7: A Success Page provides feedback to the User that the form Submission was successful.

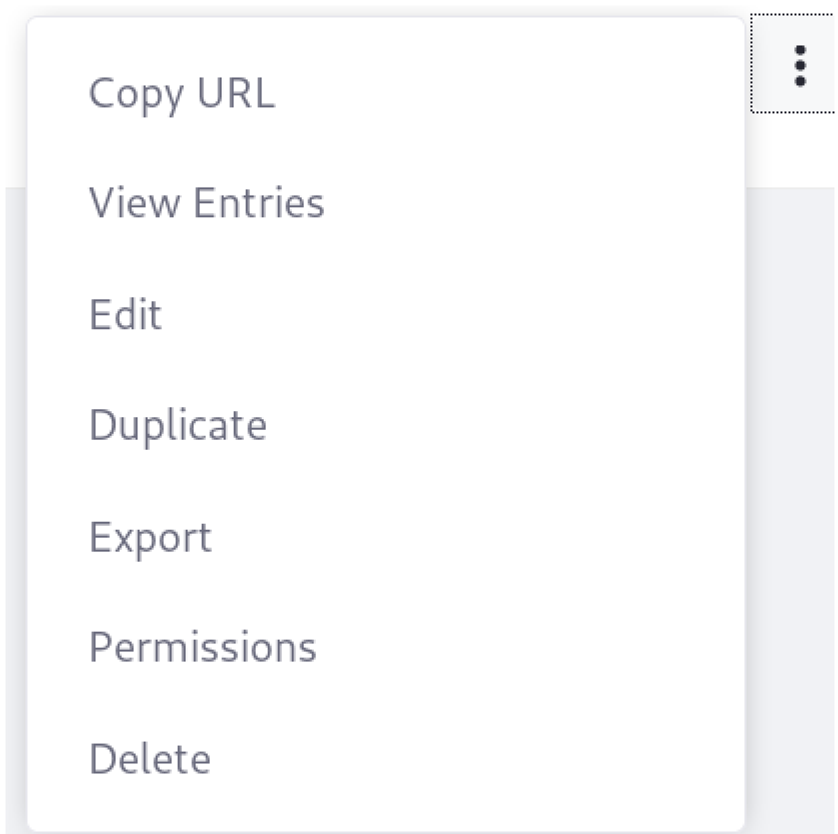

Figure 59.8: Duplicate a Form to get a head start on your next, similar Form.

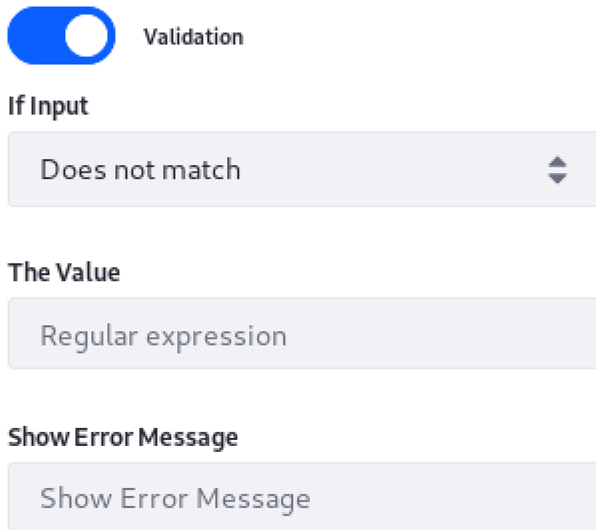

Figure 59.9: Validate Text fields to ensure Users enter the correct data.

Data Provider Enhancements Data Providers are still used to populate Select from List fields. The Data Provider setup form has been overhauled, and you can import and export Data Provider definitions. Data Providers are also linked with the Autocomplete feature and the Autofill rule.

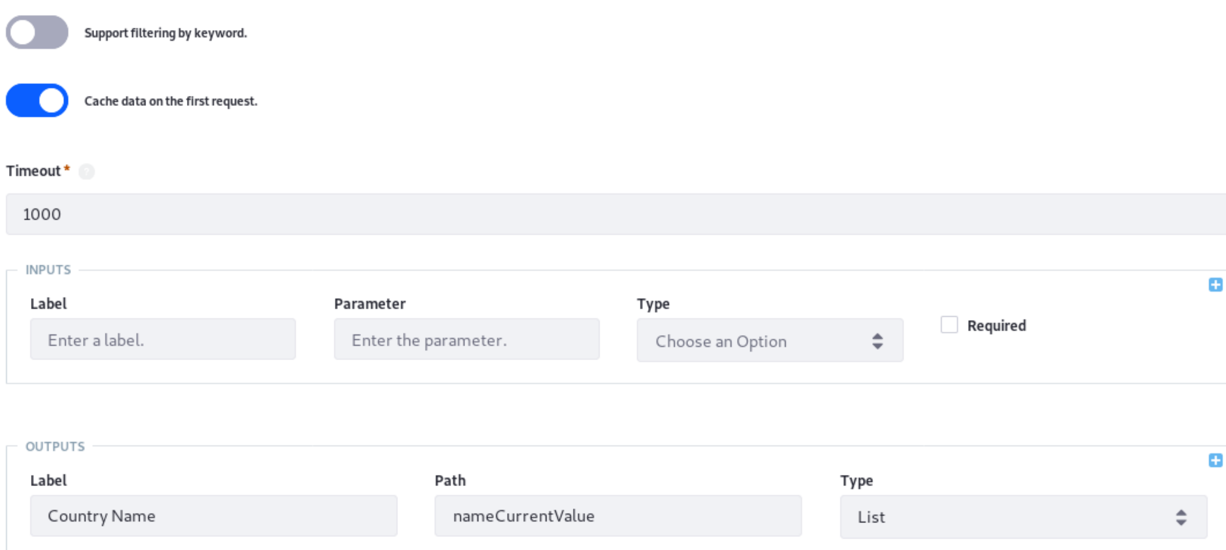

Figure 59.10: : The Data Provider interface was redesigned and new options added.

XLS Form Entries Data Export entries to XLS format.

More enhancements are planned for the future.

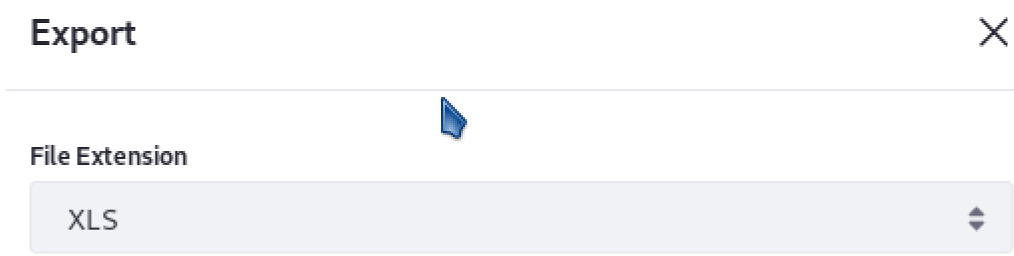

Figure 59.11: : Export Form Entries to XLS. CSV and XML are also available.

59.4 Creating and Managing Forms

The Forms widget can do a lot of things really well, but if you just need a simple form, how do you wade through all the features you don't need? Is your simple survey going to make you late for that lunch outing you've been planning with colleagues at that new shawarma place? No!

Let's create a simple yet elegant form, give access to the intended users, and get you on your way to lunch.

At The Lunar Resort, it's important to capture guests' feedback about their stay at the resort. After a (hopefully) safe journey home, guests should receive an email with a link to brief survey that prompts them to rate their stay from a list of selections, and add any additional comments they'd like in an optional field.

Viewing Forms

Whether creating a form or managing existing forms, it all starts in the same place: the Forms Application in your site's Content section. Access this in the Menu, first choosing the site to work in (for example, The Lunar Resort) and clicking *Content* → *Forms*. The first thing you'll see is a list of existing forms (if there are any). This list is styled by the Display Style selector next to the Add button (). By default, forms are displayed in List format.

There's also a Table format. Change the style for a single site right here in the Forms site menu application, or change the default display style for the system scope in Control Panel → Configuration → System Settings → Forms (in the Content section). Click the *Forms* entry and find the Default Display View property. Click *Update* and your changes are propagated to all sites.

Building a Form

To add a new form,

1. Click the *Add* button (). The form builder view appears.

2. Name the form. Replace *Untitled Form* with *Guest Survey*.

3. For Help Text enter *Tell us how your stay was!*

4. Add the fields. Click the *Add* button () to open the sidebar if it's not already opened.

5. Drag a *Select from List* field onto the form builder and configure it like this:

   Label: *Rate your visit to The Lunar Resort.*

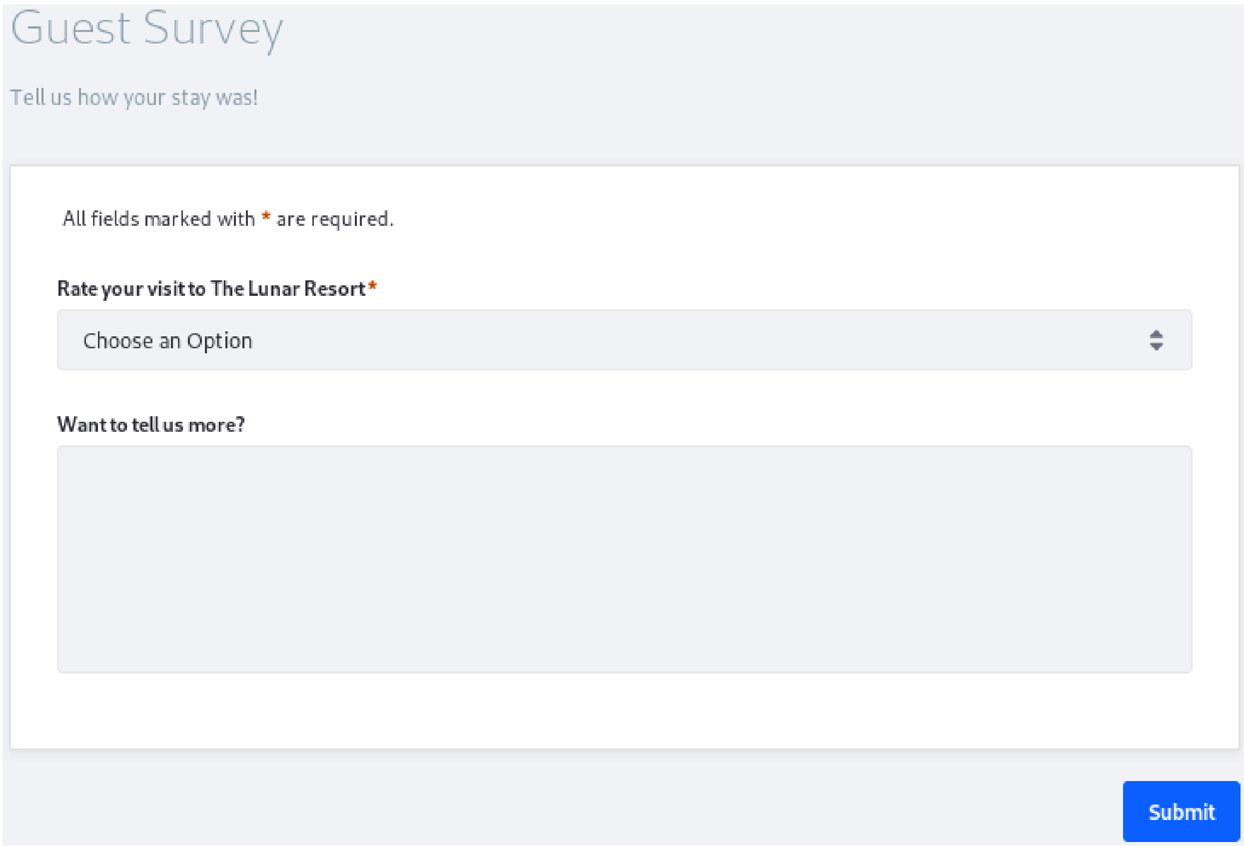

Figure 59.12: Get feedback from guests of The Lunar Resort.

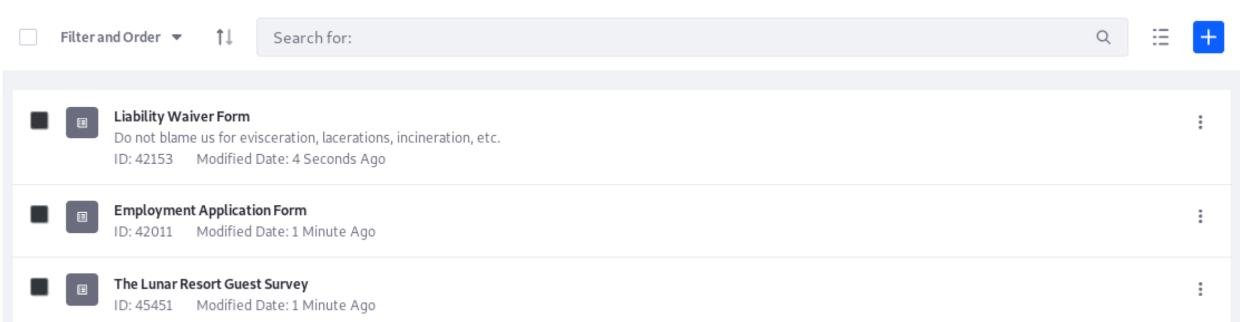

Figure 59.13: Forms are displayed in List format by default.

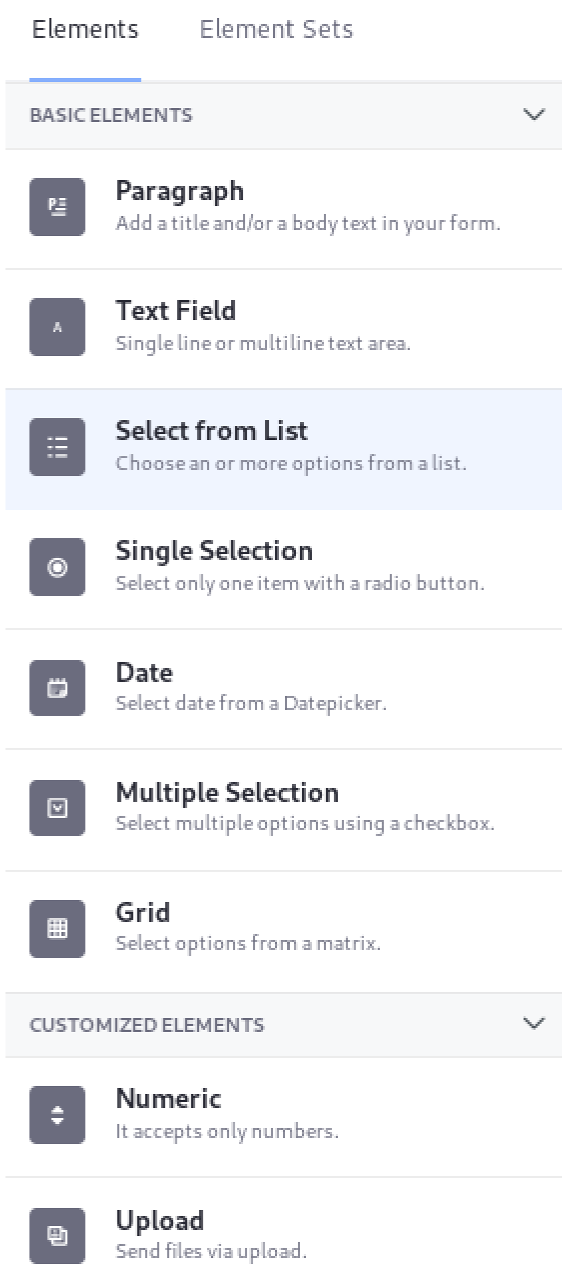

Figure 59.14: You can choose from nine field types when creating forms.

Help Text: Leave this blank for now. If you want a subheading for your field to provide additional guidance, this would be useful.

Turn on the *Required Field* selector. At a minimum, this form must capture whether guests like their stay or not.

Leave the manual option checked for creating the list of selections. To learn about populating the field with a data provider, read here.

Add these options:

- *It was out of this world!*
- *I had a good time.*
- *I'd rather go to the beach*
- *I'll never come back*

Typing in one of the fields automatically adds another blank selection line. Just leave the last one blank when you're done.

6. To see additional options, click *Properties*. Close the sidebar when finished.
7. Add a text field, using the same procedure you used for the select field.

   Label: *Want to tell us more?*

   Help Text: Leave this blank again to give the form a consistent look.

   My text field has: Choose *Multiple Lines*. Let guests prattle on about their stay if they want to.

   Required Field: Leave this unselected. Only require guests to fill out the select field and leave this one as optional.
8. Close the sidebar.
9. In the form builder, you can see the way the fields are laid out on the form page.
10. When the form is finished, click *Save Form*. It's also auto-saved every minute by default.
11. Click *Publish Form*. A dedicated URL to the form is generated, but nobody has the URL yet.

Now your form can be added to a page, and Lunar Resort guests can be emailed and provided with a link to the page where the form is displayed.

Accessing Forms

Once the form is developed and published, there are two options for getting the published form to targeted users:

1. Place the Form widget on a site page. This approach lets users navigate to the page in the site.
2. Copy the dedicated form URL and provide it to users (for example, via email). This limits access to the form to only those users who have the direct link.

To display the form on a site page in The Lunar Resort site:

1. Add a page to the site (choose Full Page Application for the page type if you only want the form application on the page). Call it *Guest Survey*.

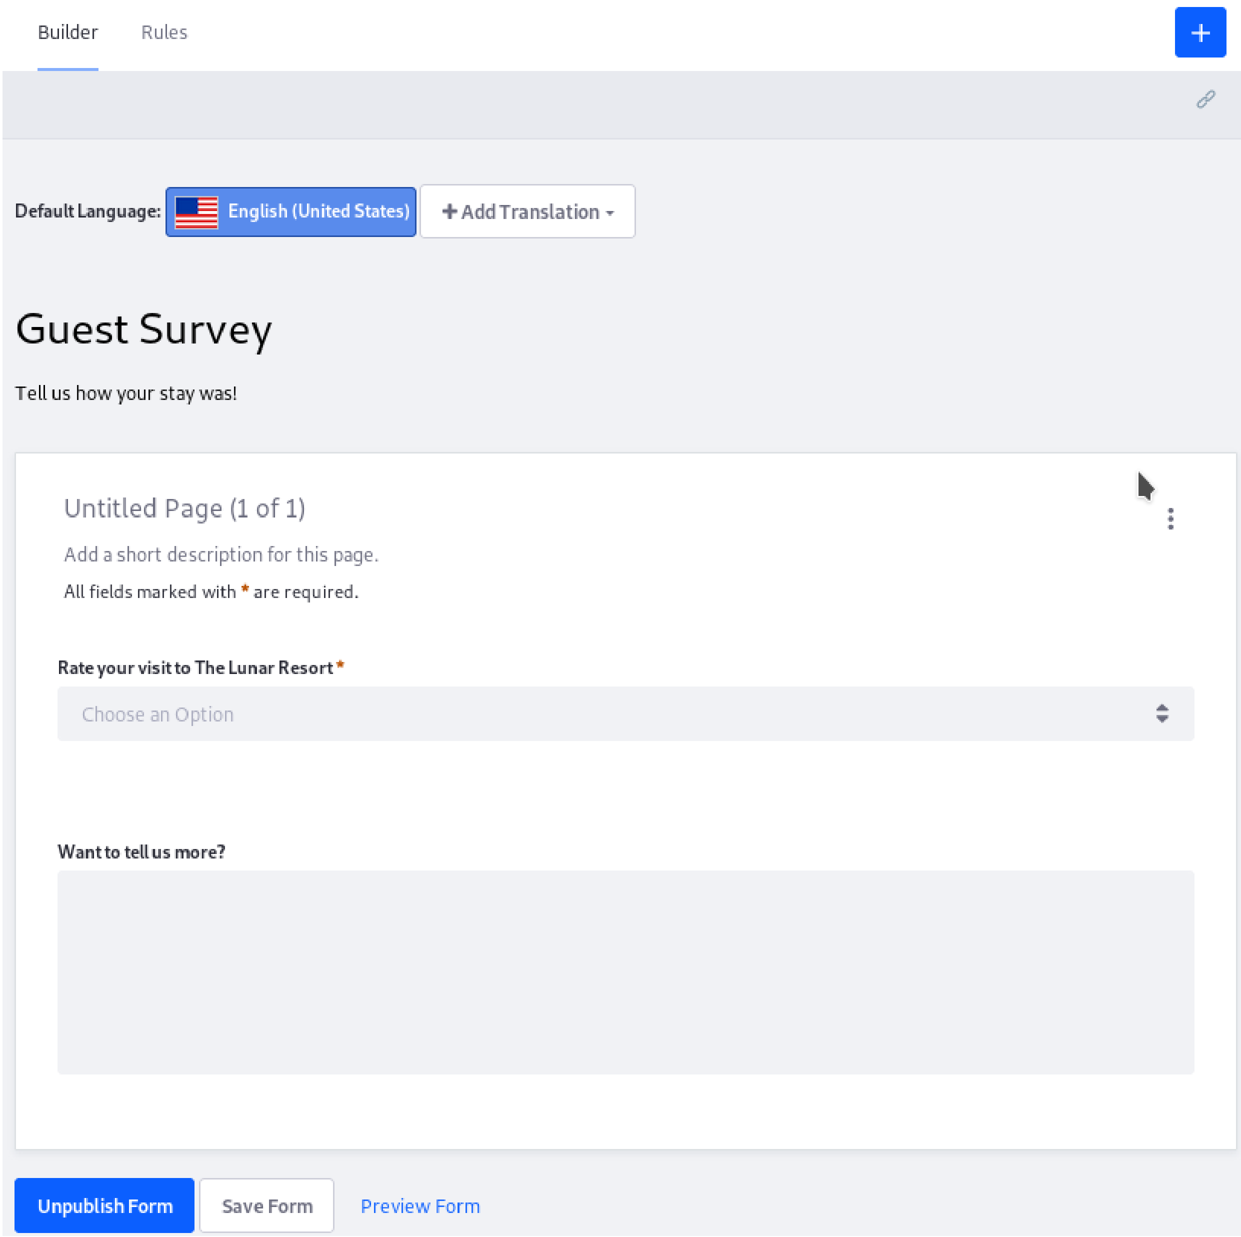

Figure 59.15: The form builder page lets you preview your form layout, add a page to the form, or add some more fields.

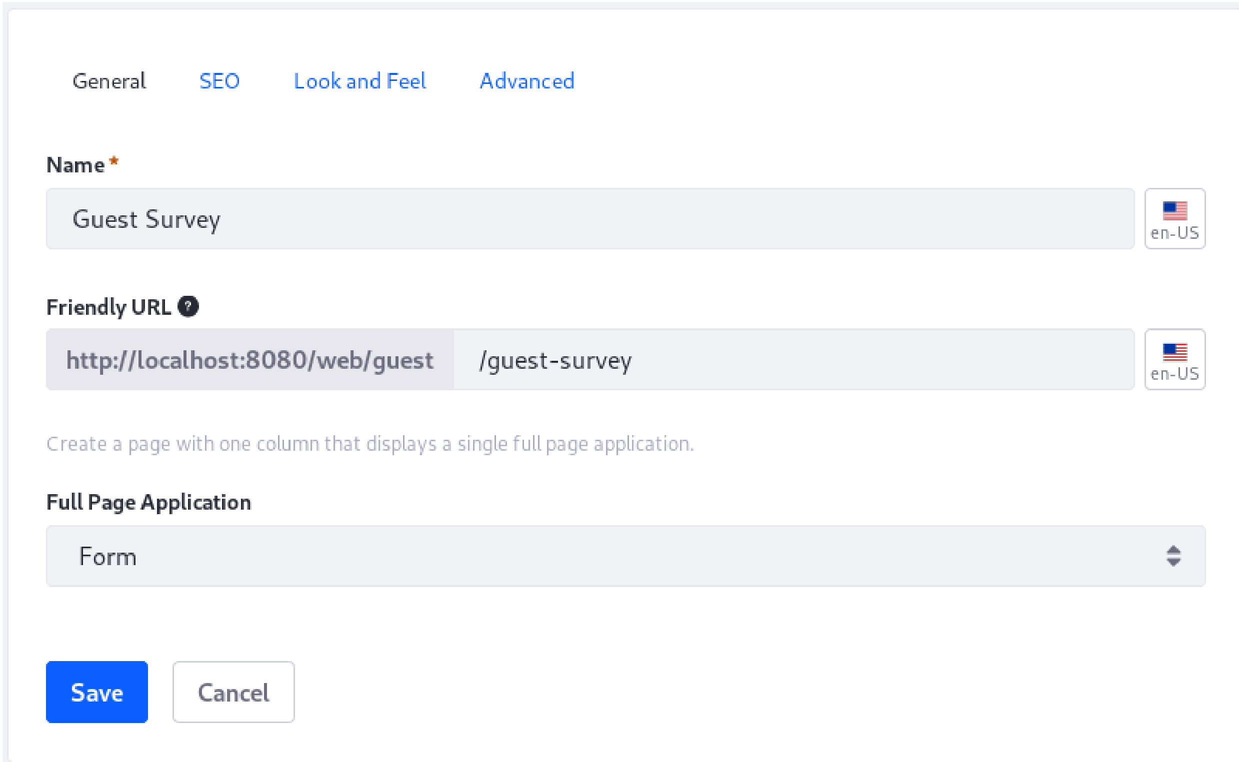

Figure 59.16: Add a page for guests to view and fill out your new form.

2. Add the Form widget to the page. If you used a full page application, click the *Configure Page* button ( ) and choose *Form* from the Full Page Application dropdown.
3. Once the Form widget is on the page, click *Select Form*, choose the *Guest Survey* form, and click *Save*.
4. Close the Form dialog window and your form is ready for Lunar Resort site users.

To display the form on a dedicated page accessed only by its link:

1. In the form builder, click *Publish Form* if you haven't already.

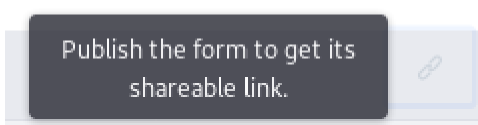

Figure 59.17: You must first publish a form before you can get a shareable link.

2. Once published, click the link icon at the top right of the builder.
3. Once you get the link out to users, it's showtime.

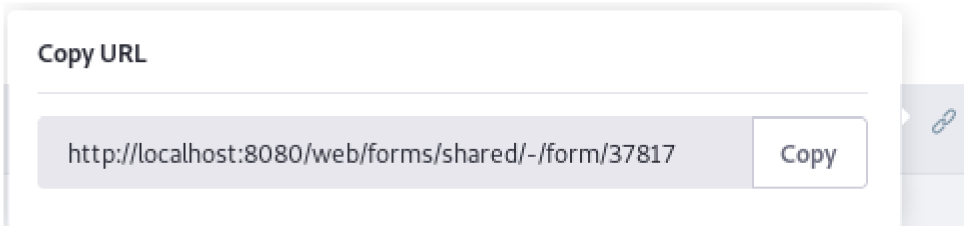

Figure 59.18: Copy the link to your form.

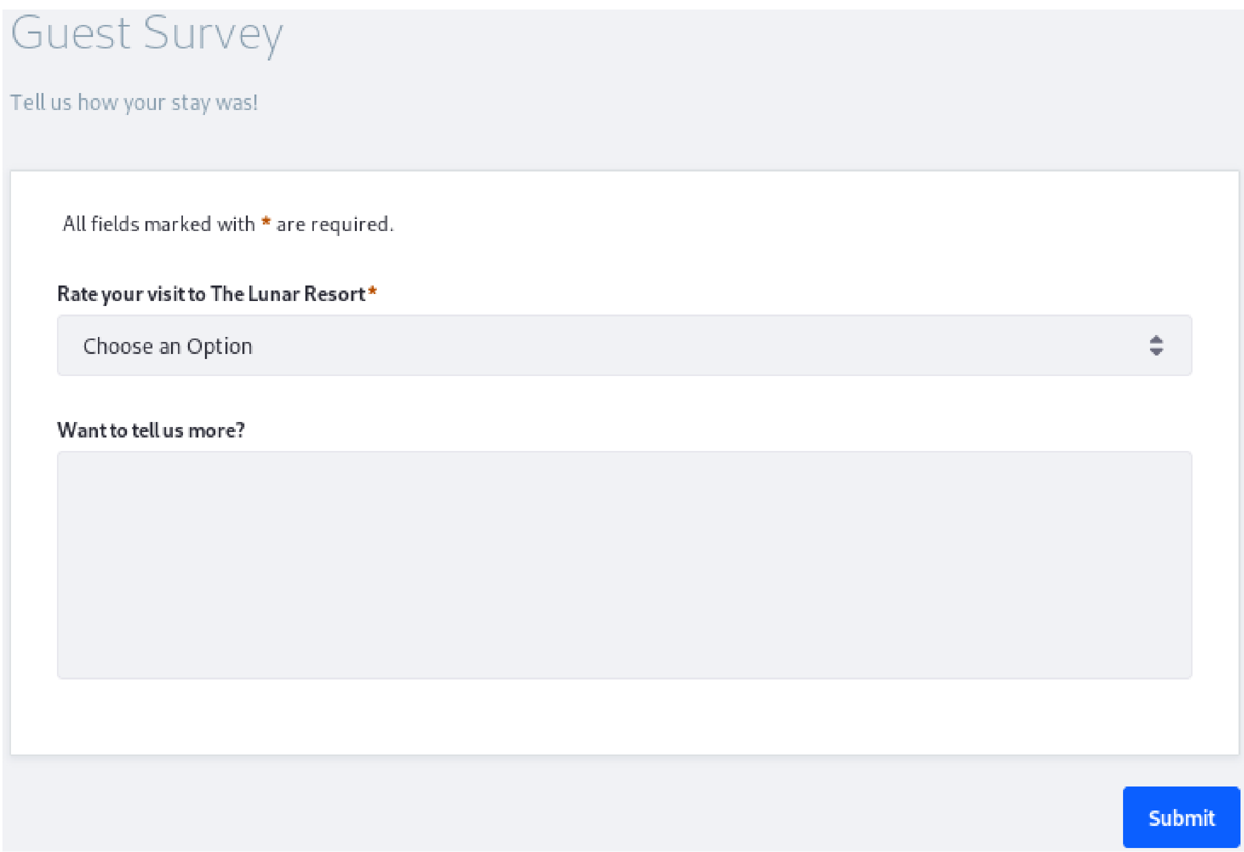

Figure 59.19: Lunar Resort guests can use a simple form to record their feelings about the resort.

Next you'll learn how to view the form entries. Since there aren't any yet, fill out and submit the form a few times.

Now you know the basics of creating and managing forms, but this presentation didn't do the Forms application justice. It's much more powerful than hinted at here. The remaining articles in this section immerse you in more advanced form building features.

59.5 Managing Form Entries

Once users begin submitting form entries, you'll want to do these things with them:

- View form entries
- Export form entries
- Delete form entries

Start by learning how to access and view the entries.

Viewing Form Entries

When users fill out forms, they're generating data. You'll want to see that data at some point.

1. From the Menu, navigate back to the *Content* → *Forms* section of The Lunar Resort site.

2. Click the *Actions* (⋮) button for the form and select *View Entries*.

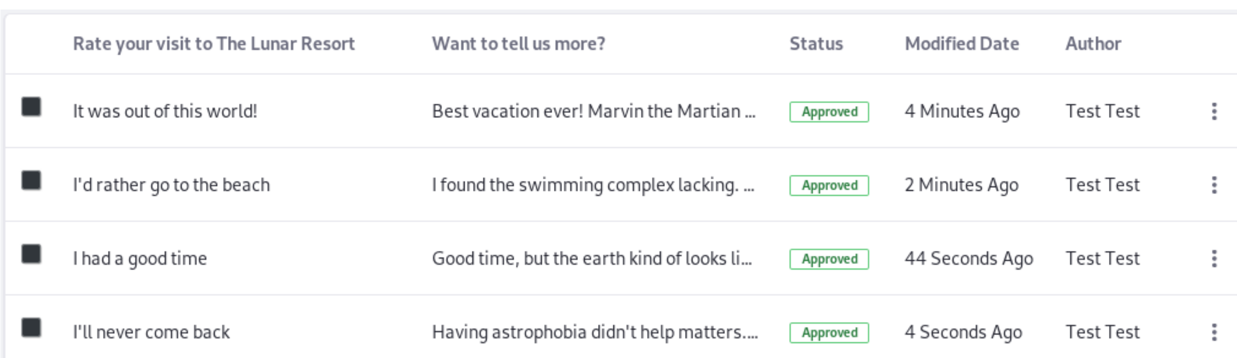

	Rate your visit to The Lunar Resort	Want to tell us more?	Status	Modified Date	Author	
	It was out of this world!	Best vacation ever! Marvin the Martian ...	Approved	4 Minutes Ago	Test Test	⋮
	I'd rather go to the beach	I found the swimming complex lacking. ...	Approved	2 Minutes Ago	Test Test	⋮
	I had a good time	Good time, but the earth kind of looks li...	Approved	44 Seconds Ago	Test Test	⋮
	I'll never come back	Having astrophobia didn't help matters....	Approved	4 Seconds Ago	Test Test	⋮

Figure 59.20: You can view the entries right in the Forms application.

3. What if you have a lot of form fields and you can't see all the data for each entry in the search container? Just click the *Actions* (⋮) button for the entry and select *View*. You're shown all the specifics for that form entry.

Viewing entries is great, but this is serious data we're talking about. You might need to get all the entries into a spreadsheet so you can work with them.

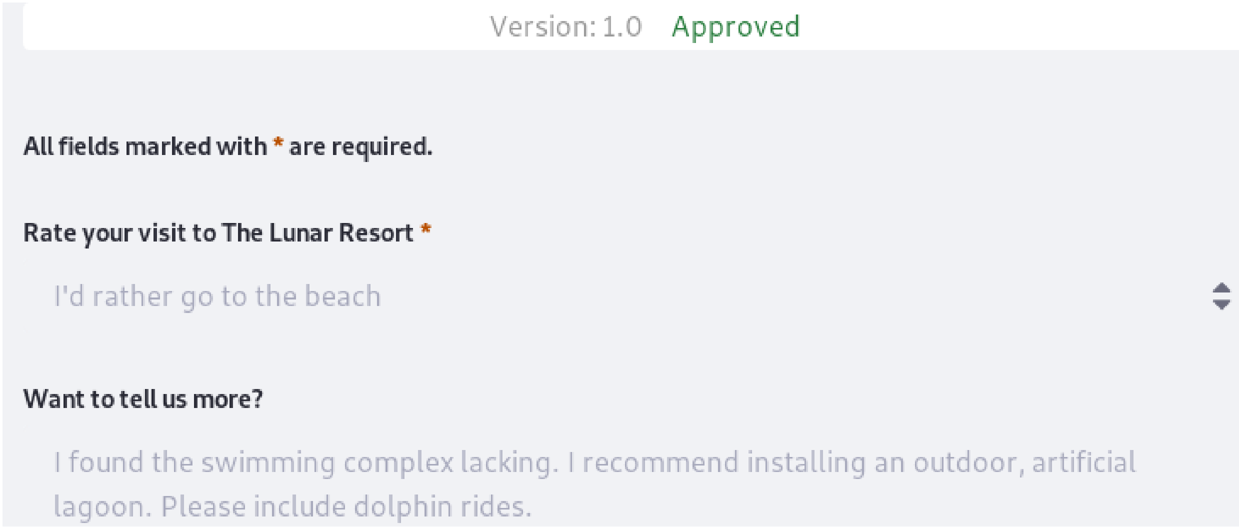

Figure 59.21: You can view a single entry right in the Forms application.

Exporting Form Entries

So you need to put your form entries in a spreadsheet to do spreadsheet things with them? No problem.

1. Navigate to the Forms application in The Lunar Resort site's Content section again.
2. Click the *Actions* ( ) button and select *Export*.
3. Choose a File Extension. You can export entries in CSV, XLS, or XML formats by default. For this example, pick CSV.
4. Click *Okay*, and open the file or save it locally. Open it with your favorite spreadsheet program and verify your form entries.

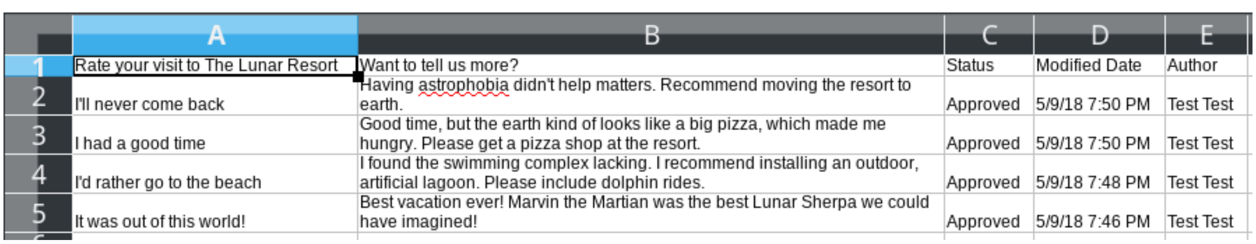

	A	B	C	D	E
1	Rate your visit to The Lunar Resort	Want to tell us more?	Status	Modified Date	Author
2	I'll never come back	Having astrophobia didn't help matters. Recommend moving the resort to earth.	Approved	5/9/18 7:50 PM	Test Test
3	I had a good time	Good time, but the earth kind of looks like a big pizza, which made me hungry. Please get a pizza shop at the resort.	Approved	5/9/18 7:50 PM	Test Test
4	I'd rather go to the beach	I found the swimming complex lacking. I recommend installing an outdoor, artificial lagoon. Please include dolphin rides.	Approved	5/9/18 7:48 PM	Test Test
5	It was out of this world!	Best vacation ever! Marvin the Martian was the best Lunar Sherpa we could have imagined!	Approved	5/9/18 7:46 PM	Test Test

Figure 59.22: You can export entries as CSV, XLS, or XML.

Note: The Forms application itself has an *Import/Export* window accessible from the application's Configuration menu ( ). This is how you import and export the application configuration and its data (forms and form entries). The file format for this type of import and export is a LAR file. For more information, see the article on importing and exporting application content.

There's a system level setting to determine whether administrators can export entries in CSV format:

1. Go to Control Panel → Configuration → System Settings and click the *Forms* category in the Content section.

2. Click the *Forms* entry under SITE SCOPE.

3. The CSV Export property has three options:

    - *Enabled* to enable CSV Export without a warning
    - *Enabled (Show Warning)* to enable CSV Export with the following warning to administrators: This CSV file contains user supplied inputs. Opening a CSV file in a spreadsheet program may be dangerous.
    - *Disabled* to turn off CSV Export.

Once you export a batch of form entries, it can make sense to delete them from the database.

Deleting Form Entries

What if you export a form's entries and now you want to remove them from the Liferay database? It's easy to delete all of a form's entries at once.

1. Navigate back to the Forms application In The Lunar Resort Content section.

2. Click the *Actions* ( ) button next to the Guest Survey form and select *View Entries* again.

3. Select all entries by checking the box next to *Filter and Order*. A trash can icon ( ) appears in the top right corner of the Form Entries screen. Click it.

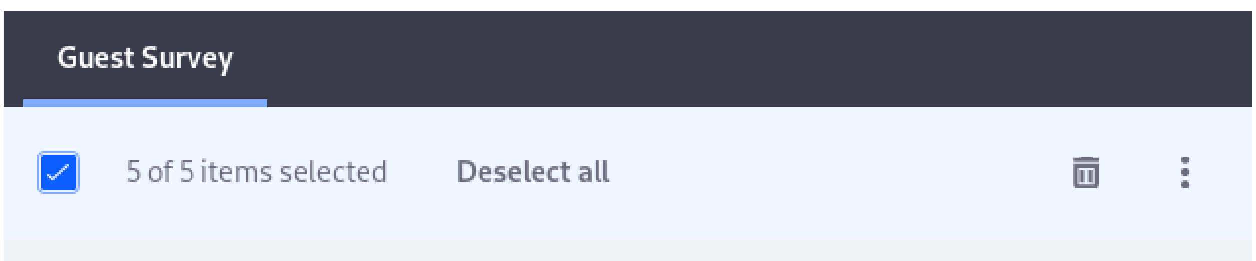

Figure 59.23: Delete all form entries in one fell swoop.

If you just wanted to delete a single entry, select that entry by checking its box; then click the trash can.

If you're worried about deleting everything irrecoverably by accident, don't worry. You must confirm the deletion in a dialog box that pops up after clicking the trash can.

Now you can create basic forms and manage the entries. Keep reading in this section to learn about the many additional form building features available to you.

Chapter 60

Form Rules

What's the difference between a chicken and a dog? Among other things, you can't train a chicken, while dogs are quite trainable. If you're skeptical, try teaching your chicken to sit on command or herd sheep. Better yet, get a team of chickens to pull a sled in the Iditarod. The Forms application is much more like the dog than the useful (southwestern omelet anyone?) but untrainable chicken, and it's only getting more trainable as time passes.

Form rules are a good example of the trainable nature of the Forms application. With form rules, you can train your form fields to behave as you wish. There are several things you can make them do:

Show/hide Based on a predefined condition, set the visibility of a form field.
Enable/disable Use a predefined condition to enable or disable a field.
Require Use a predefined condition to make a field required.
Jump to Page Based on user input, skip over some form pages directly to a relevant page.
Autofill with Data Provider Use a data provider to populate fields when a condition is met in another field.
Calculate Populate a field with a calculated value using data entered in other fields.

Form rules are for changing fields and form elements by acting on conditions.

If [condition] do [action].

If you're not already familiar with the Forms application, start here. Once you know how to create forms, add and configure fields, and then publish forms, come back here and learn about form rules.

60.1 The Anatomy of a Form Rule

Each rule consists of one or more conditions and actions.

Conditions determine whether any actions are executed.

Actions determine what happens if the condition is met.

Rules are stored in the database in JSON format by default.

60.2 Creating Form Rules: Rule Builder

Once you create a form and lay out its fields, you're ready to set up rules in your form:

1. Save the form.
2. Open the Rule Builder by clicking the *Rules* tab at the top of the *Edit Form* screen.
3. In the rule builder view, you can now begin developing your form rule. Click the *Add* button (+) to get started.

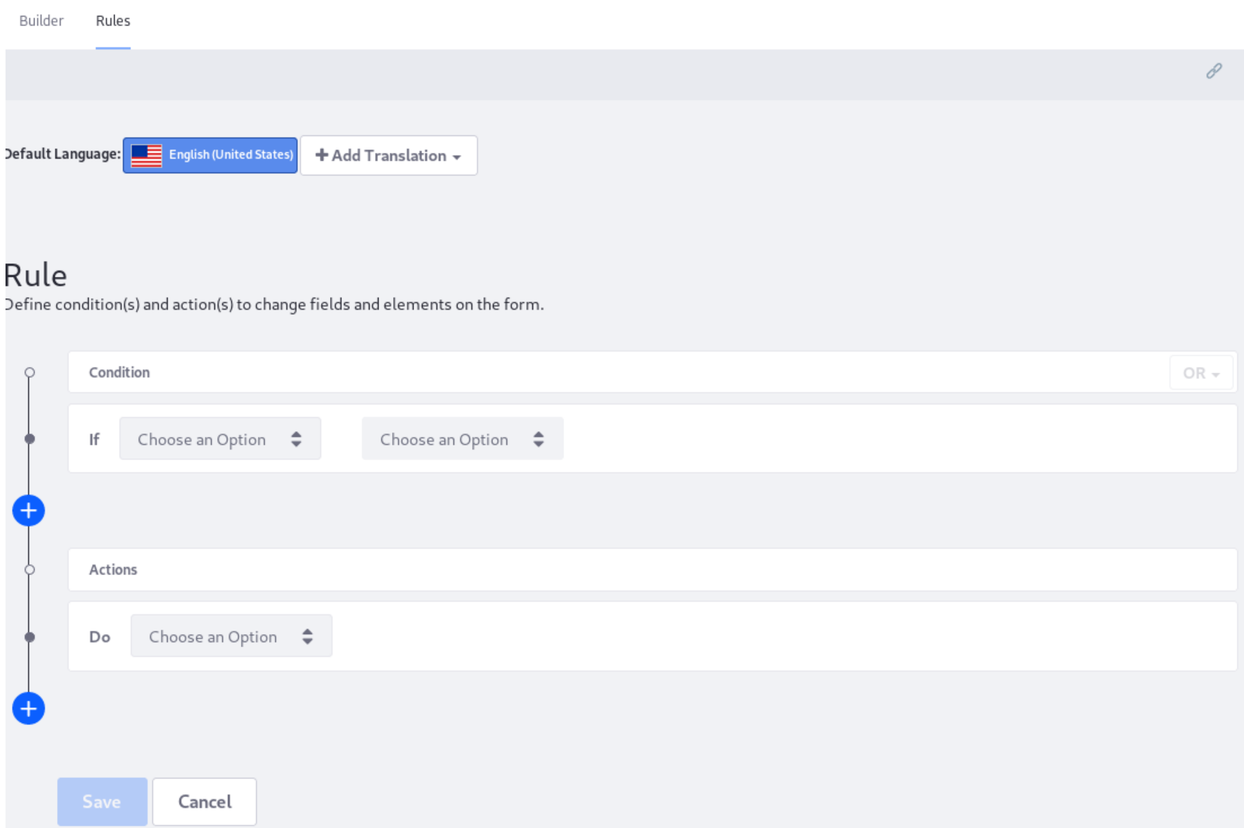

Figure 60.1: The Rule Builder gives you a handy interface for creating dynamic form rules.

Before looking at each type of rule condition and action you can use to develop rules, consider the *OR* selector box at the right side of the *Condition* (it's grayed out and unusable at first). You can choose *OR* or *AND* here, depending on what relationship the conditions should have with the action.

OR The action is triggered if *any* of the conditions you specify evaluates to *true*
AND The action is triggered only if *all* the conditions you specify evaluate to *true*

This box becomes usable once you click the Add button (+) to add an extra condition.

60.3 Conditions

Conditions are the gatekeepers of form rules. If the condition's *if statement* evaluates to *true*, the action is triggered. If it evaluates to *false*, no action happens.

A condition checks whether one field's value

- *Is equal to* a specific value or another field's value.
- *Is not equal to* a specific value or another field's value.
- *Contains* a specific value or another field's value.
- *Does not contain* a specific value or another field's value.
- *Is empty*. This assumes you want to do something if a field *is* empty.
- *Is not empty*. This assumes you want to do something as long as a field is *not* empty.

One exception to this is the User condition, which is the last option in the Condition dropdown menu.

The User condition doesn't act on a field at all. It checks whether a User belongs to a certain role. For example, if the condition

If `User` *belongs to* `Administrator`

evaluates to *true*, an action is triggered.

A condition is the gateway into a form rule, but actions define what actually happens when the condition evaluates to *true*. The remaining articles discuss the available actions and demonstrate their use.

60.4 Action: Show and Hide

With a show and hide rule, you use one or more conditions to determine whether to show or hide a field if the condition evaluates to *true*.

To set this example up, add these fields to a form:

- *I am 18 Years Old or Older*, a required single selection field with two options: *Yes* and *No*.
- *Legal Guardian Email Address*, a text field that accepts valid email addresses (use text field validation to dictate input type).

Example: If you're under 18 years old, you need the approval of a legal guardian to drive a sled in a sled dog race (even if you're racing chickens, not dogs). The form for registering your chicken team asks you the age of the driver. If you enter a number less than 18, the Legal Guardian Email Address field appears.

To configure a Show/Hide rule,

1. Open the Rules tab of the Edit Form page and click the Add () button.
2. Define the rule:
   - If the *I am 18 years old or older* field is equal to *No*, show the *Legal Guardian Email Address* field.

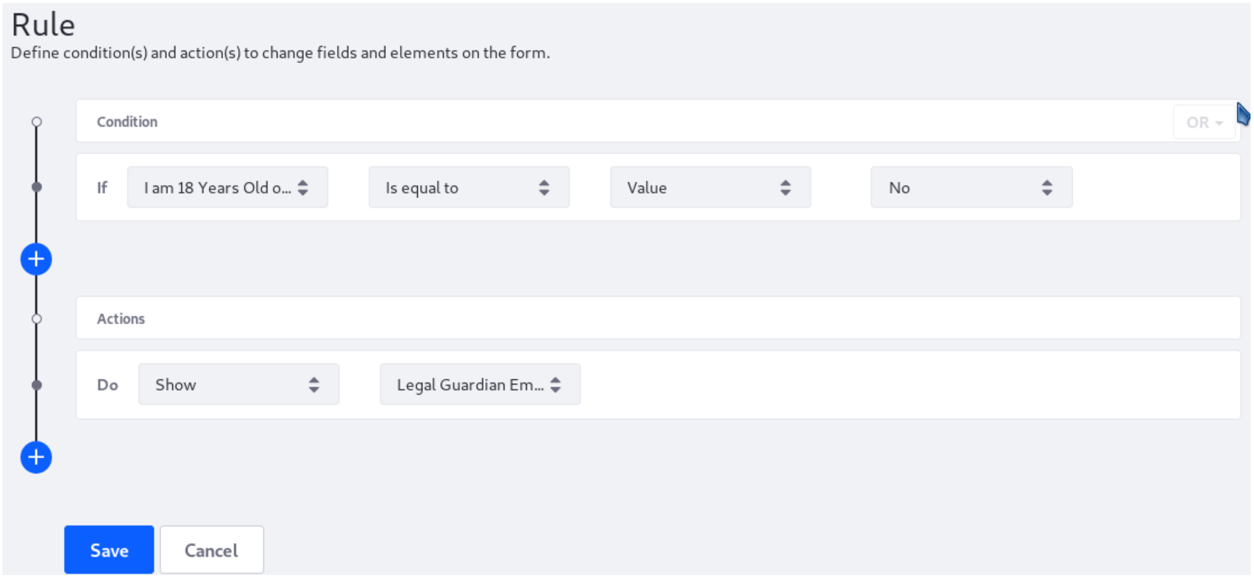

Figure 60.2: Build form rules quickly by defining your conditions and actions.

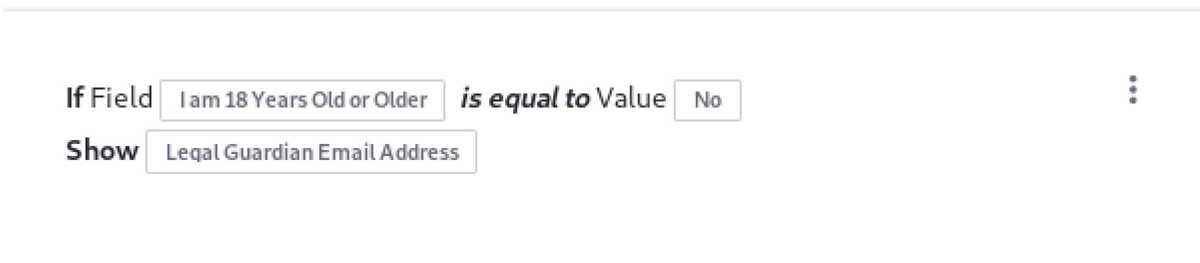

Figure 60.3: Once a rule is saved, it is displayed so that you can easily understand what it does.

- Save the rule.

Now the *Legal Guardian Email Address* field is only displayed in the form if the user selects *No* in the *I am 18 years old or older* field.

Show rules let you keep a field hidden until some condition is met.

60.5 Action: Require

Use a require rule to make a field required based on one or more conditions.

Example: If you are following the example, you already set up a *show* rule, where a *Legal Guardian Email Address* field appears if the user selects *No* in the *I am 18 years old or older* field. You also want to make the *Legal Guardian Email Address* field required.

To configure a require rule,

1. Edit the Show Rule configured above. Open the Rules tab of the Edit Form page and click the kebab menu (⋮) for the rule, and then click *Edit*.
2. Add an Action to the rule:

   - If the *I am 18 years old or older* field is equal to *No*, show the *Legal Guardian Email Address* field and make it required.

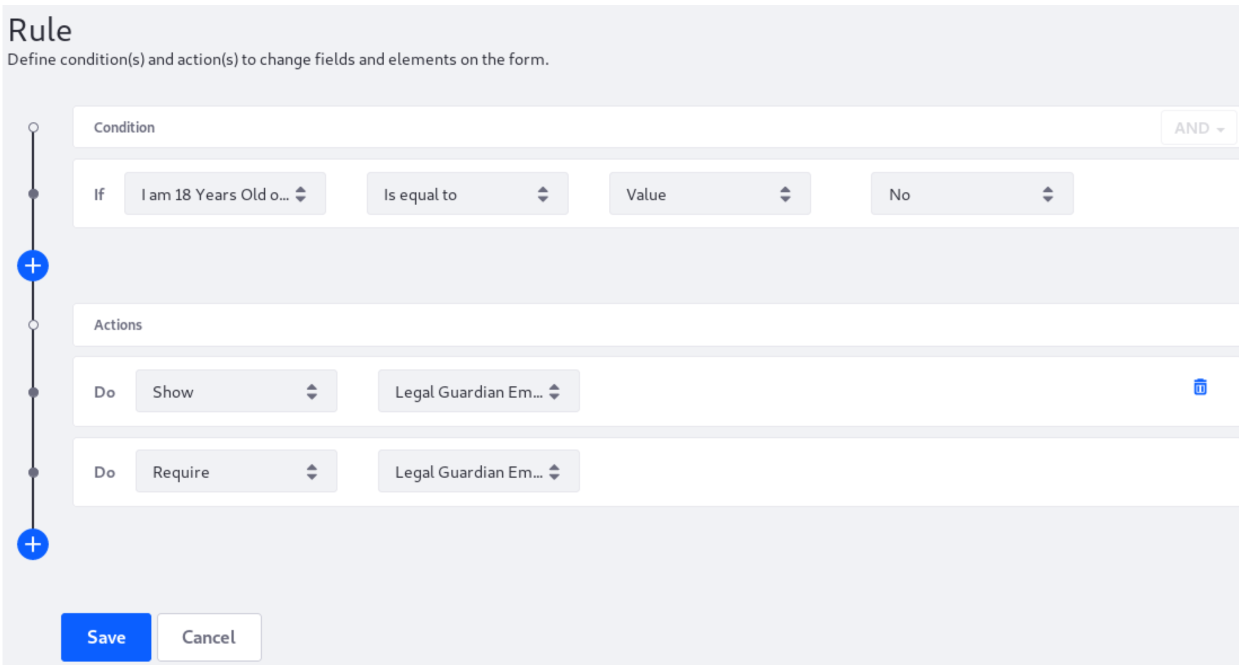

Figure 60.4: Build form rules quickly by defining your conditions and actions.

   - Save the rule.

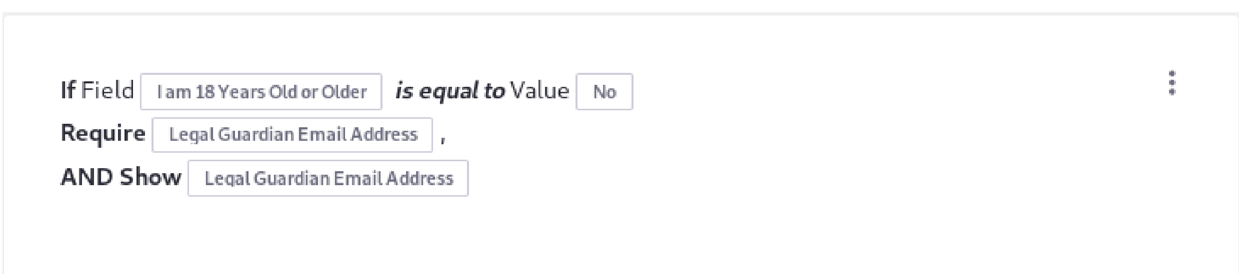

Figure 60.5: Once a rule is saved, it is displayed so that you can easily understand what it does.

Require rules let you require fields based on input from other fields.

60.6 Action: Enable and Disable

Use an enable/disable rule to make a field editable based on one or more conditions.

Example: Part of the race registration fee pays for dog food. You don't have to feed your chicken team with the provided dog food though. There's a single selection field that asks *Would you like use the provided dog food?*. If you select *Yes*, you can select how much food, in US pounds, you'll need for your team throughout the race. Since you're racing chickens, you'll select *No*, and the *Amount (US lb.)* field is disabled.

To follow the example, add a single selection field *Would you like to use the provided dog food?* with two options: *Yes* and *No*.

Add a numeric field called *Amount (US lb.)* and make it an Integer. Use field validation to make sure it's not greater than *100*.

To set up the enable/disable rule,

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:

   - If the *Would you like to use the provided dog food?* field is equal to *Yes*, enable the *Amount (US lb.)* field.

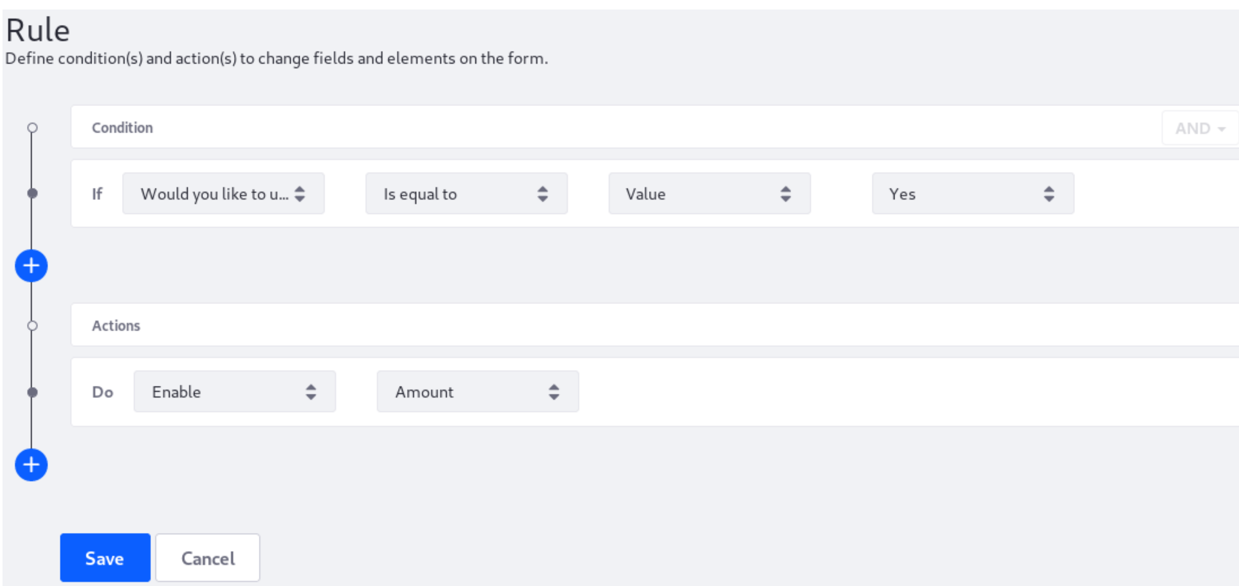

Figure 60.6: Build form rules quickly by defining your conditions and actions.

   - Save the rule.

Now users can fill out the amount of dog food they'll need only if they specify that they do indeed want to use the provided food.

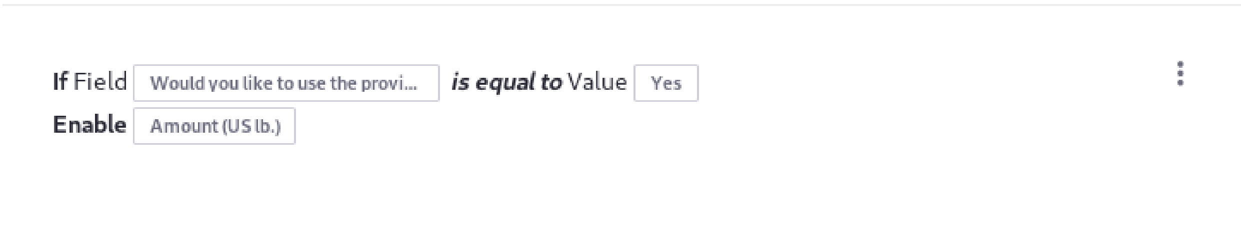

Figure 60.7: Once a rule is saved, it is displayed so that you can easily understand what it does.

60.7 Action: Jump to Page

Use a Jump to Page rule to navigate automatically to a specific page in the form based on one or more conditions. This is useful if some pages don't apply to all the form's users. Even fields marked as required on the skipped pages can be successfully skipped using this rule.

This action doesn't appear in the rule builder unless the form has multiple pages.

To follow the example here,

1. Create a second form page called *Team Information*.
2. On the new form page, create a single selection field labeled *Are you a returning racer with the same team?* with two options: *Yes* and *No*.
3. On the new form page, create a repeatable text field called *Animal Name*.
4. Create a third form page called *Final Confirmation*.

Example: There's a question on the *Team Information* page of the dog sled race registration form that asks *Are you a returning racer with the same team?* If you select *Yes*, when you click the form's *Next* button, you skip to the final page of the form, since there's no need to fill out your animal's name again. Their monogrammed T-shirts will be ready at the start of the race.

To configure the Jump to Page rule,

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:
    - If field *Are you a returning racer with the same team?* is equal to *Yes*, Jump to Page *Final Confirmation*.
    - Save the rule.

If you use an *is not equal to* condition for form fields on two different pages, the condition is checked after leaving the page of the first form field, and evaluates to *true* since there's a value in the first field and no value in the second field. It's best to use this condition with fields existing on the same page.

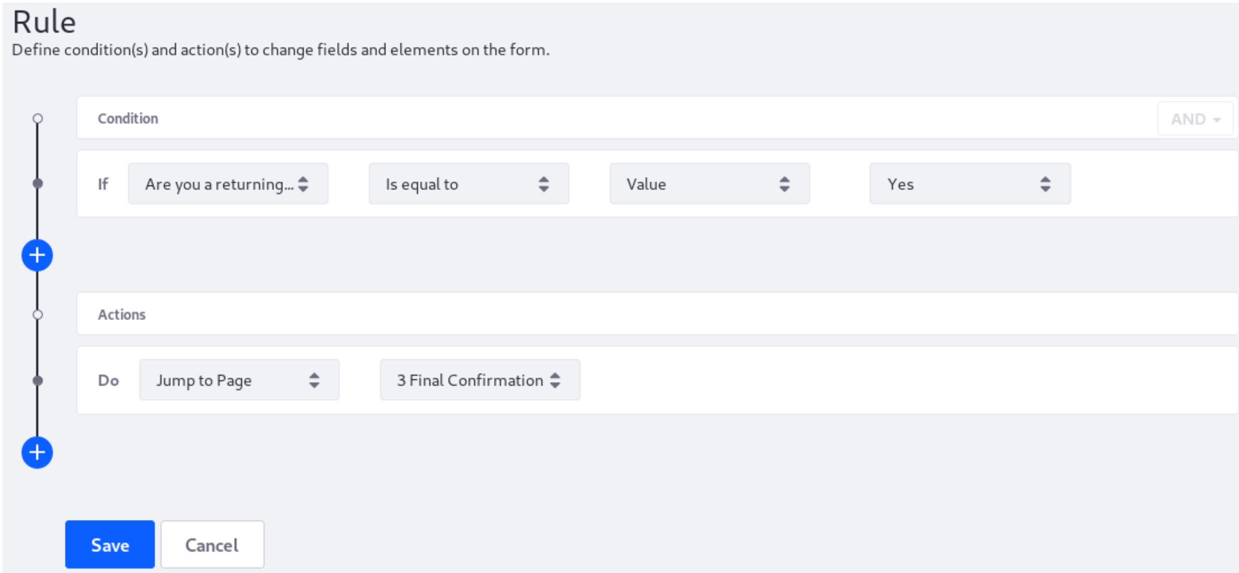

Figure 60.8: Build form rules quickly by defining your conditions and actions.

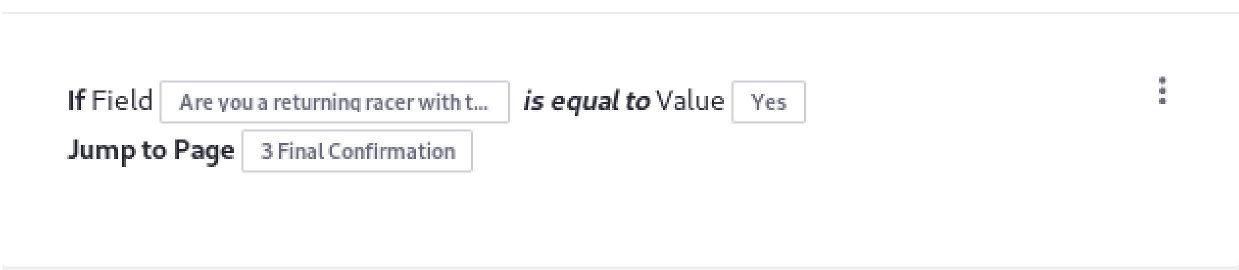

Figure 60.9: Once a rule is saved, it is displayed so that you can easily understand what it does.

60.8 Action: Autofill

Autofill rules let you change the selection options of another field based on the value entered into a related field. A data provider's output is used to populate a field, as long as the condition is met.

Before configuring an autofill rule, set up a data provider. That's how autofilled fields are populated. Pay careful attention to the input and output parameters you choose when setting up the rule.

To follow this example:

1. Set up a data provider using the `get-countries` JSON web service. If you're running Liferay DXP at `localhost:8080`, you can access this web service here:

   ```
   http://localhost:8080/api/jsonws?contextName=&signature=%2Fcountry%2Fget-countries-0
   ```

   Make sure the output parameter is set to `$..nameCurentValue`. If you're unsure how to do this, first read the article on Data Providers.

2. On the last form page, add two fields:

   - A Single Selection field called *If I win I'd like my award to be:*, with two choices: *Cash* and *All Expenses Paid Vacation*.
   - A Select from List field called *Choose a Destination Country*. Under *Create List*, choose *From Autofill*.

Example: Before submitting the race registration, let users decide whether they want a cash prize or an all-expenses-paid vacation. If they choose the vacation, populate the Choose a Destination Country with output from the data provider.

To configure an Autofill rule,

1. Open the Rules tab of the Edit Form page and click the Add () button.
2. Define the rule:

   - If field *If I win I'd like my award to be* is equal to *All Expenses Paid Vacation*, Autofill the *Choose a Destination Country* field from the *countries* data provider (note that you might have named this differently when setting it up).

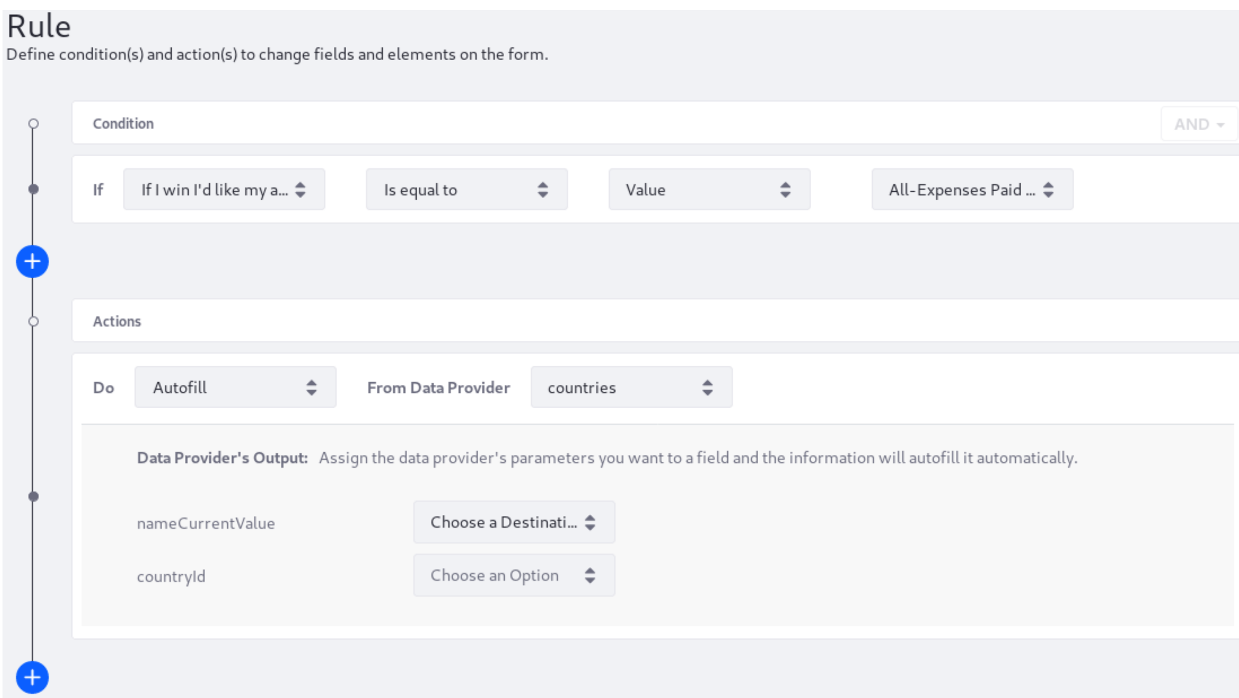

Figure 60.10: Build form rules quickly by defining your conditions and actions.

   - Save the rule.

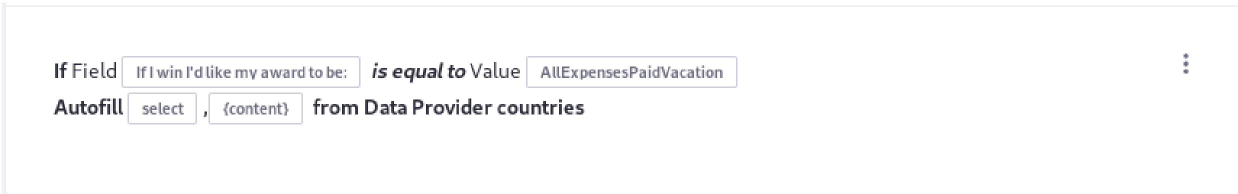

Figure 60.11: Once a rule is saved, it is displayed so that you can easily understand what it does.

Using Inputs with Autofill

The above example is simple, using only an Output to autofill a Select from List field if another field has a certain value. Many times, the response from the REST provider must be filtered before display in the Select from List field. For this, a Data Provider Input field is required. For example, to configure an autofill rule to display the countries of the world filtered by a Region field (for example, Americas, Europe, or Oceania),

1. Create a Data Provider.

   Name: `restcountries`

   URL: `https://restcountries.eu/rest/v2/region/{region}?fields=name`

   Inputs: Fill in the Label (*region*), Parameter (*region*), and Type (*Text*).

   Outputs: fill out a Label (*name*), Path (*$..name*), and Type (*List*).

   To understand more about these values, read the Data Provider documentation.

2. Create a form with these fields:

   Text: Use the Label *Region*.

   Select from List: Label it *Country*, and choose *From Autofill* under Create List.

3. Configure the Autofill rule.

   Condition: If *Region* **Is not Empty**

   Action: Do **Autofill** From Data Provider `restcountries`, Data Provider's Input: region—*Region*, Data Provider's Output: name—*Country*.

Once you're done, publish the form and try it out, by entering a valid region into the Region field, and observing that the options in the Select from List Field are filtered based on the Region. The restcountries.eu service has these regions you can use: Africa, Americas, Asia, Europe, Oceania, and Polar.

Autofill rules combine the power of data providers and form rules.

60.9 Action: Calculate

Calculate rules let you automatically populate a numeric field by calculating its value based on other fields. Calculations are limited to numeric fields.

To follow the example below:

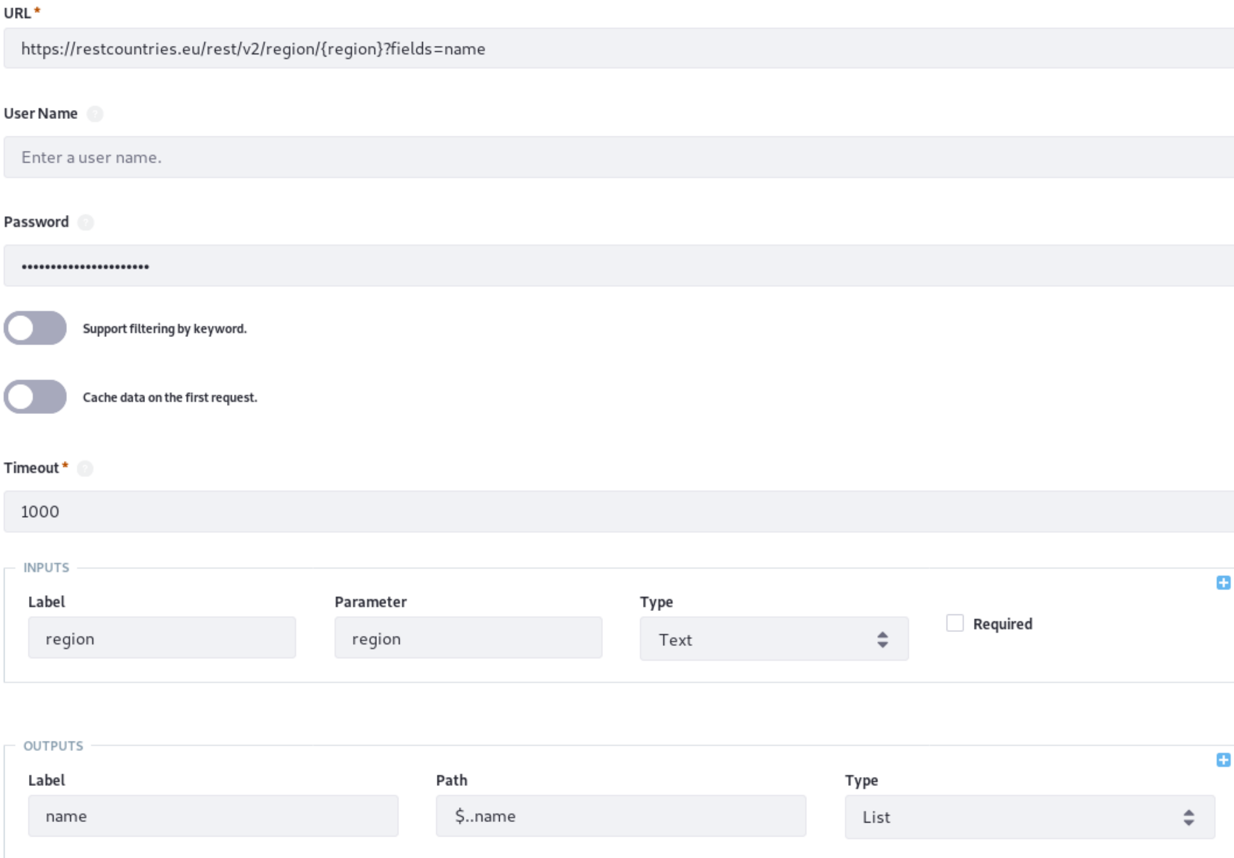

Figure 60.12: Create a data provider for the autofill rule.

1. Create 5 Numeric fields called *Animal Weight 1 (US lb.)*, *Animal Weight 2 (US lb.)*, etc.
2. Create a Numeric field called *Total Food Required (US lb.)*.

Example: A 16-dog sled team can consume 2,000 US lb. of food during the Iditarod. This equates to about 0.25 lb. of food per lb. of animal, if the race lasts ten days. We'll use five numeric fields for animal weight instead of sixteen here, because it's tedious to create sixteen fields, even with the field duplication function. When the form user enters the weight of each animal the Total Food Required field should be calculated based on this simple formula:

```
Animal Weight 1, Animal Weight 2, ... = AW1, AW2, ...
Total Food Required = TFW

(AW1 + AW2 + AW3 + AW4 + AW5) * 0.25 = TFW
```

To configure a calculate rule:

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:

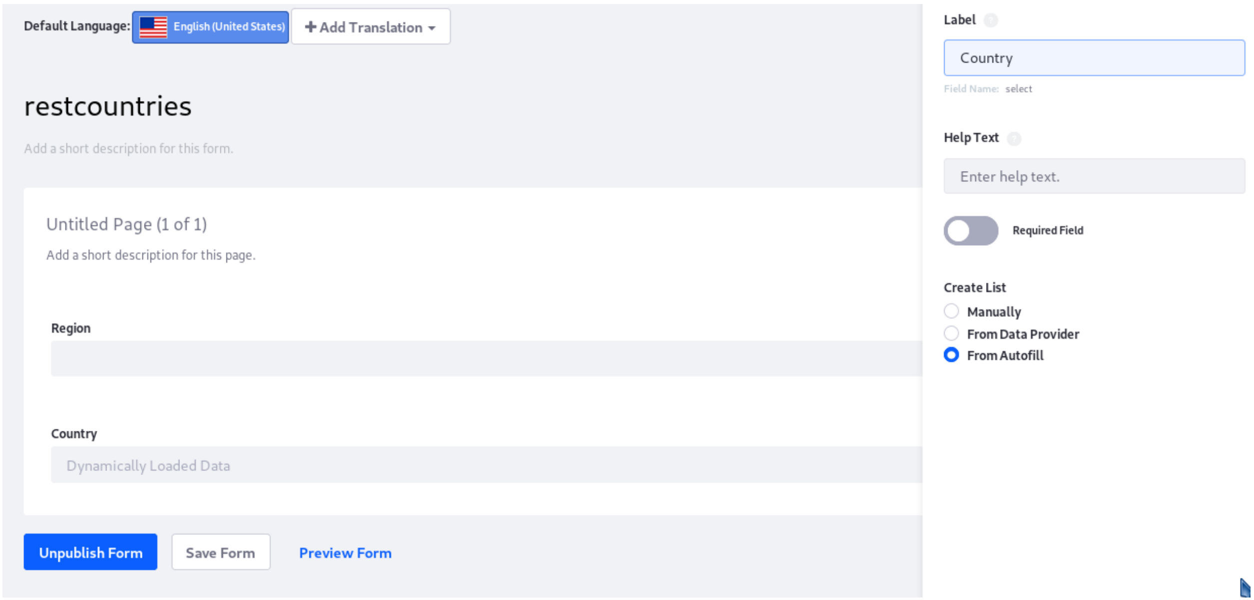

Figure 60.13: Create a form with a text field and a select from list field. These are used to provide the input to the data provider and be autofilled by its output.

- If field *Animal Weight 1* is greater than 0, Calculate the sum of the Animal Weight fields, multiplied by 0.25.
- Save the rule.

The calculation is defined using the embedded calculator. Use a mix of numeric field values, mathematical operators, and constants to define calculation rules.

60.10 Form Element Sets

If you're here looking for information on reusable field sets, you're in the right place. We call them Element Sets in the Liferay Forms application because these sets include more than just fields: they include the layout and configuration of the fields as well. In the future, additional styling elements will be available here, too.

Element sets are more like composable Form fragments or reusable Form blocks.

Sometimes you might be able to create an entire form by composing existing Element Sets. Your colleagues might call you lazy, but we'd call you industrious.

Creating Element Sets

To create Element Sets, go to Site Menu → Content → Forms. The Forms view is displayed by default. Click the *Element Sets* tab, and any existing Element Sets appear, just like existing Forms are displayed in the Forms view. Click the *Add* button ().

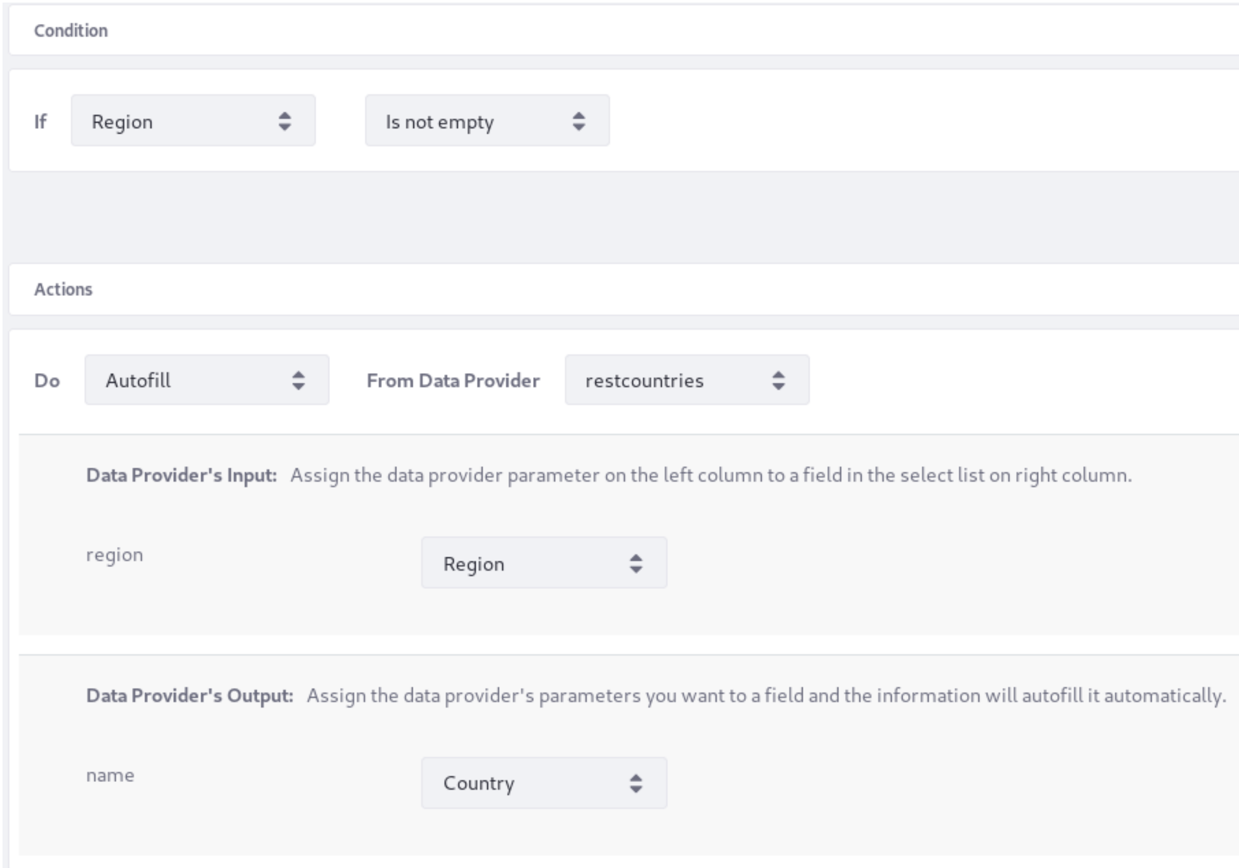

Figure 60.14: Create the autofill rule. Brag of your prowess.

Here's the thing. If you know how to create a Form, you already know how to create an Element Set. The process is identical. Drag and drop elements onto the form builder palette, configuring fields as you go.

When you're finished, click *Save*. Element Sets aren't publishable, so there's no button for that.

Once an Element Set is saved, it's instantly available for use, even in the same Element Set. That's right, you can use Element Sets to create Element Sets.

Using Element Sets

To use an Element Set in a Form:

1. Open the Form Builder.

2. If the Add Elements sidebar isn't already displayed, open it by clicking the *Add* button ( ).

3. The default view in the Add Elements sidebar is Elements. Instead click *Element Sets*.

4. Drag the Element Set onto the Form Builder, just like you would any single Form Element.

You probably already guessed that the process is identical for using Element Sets to build other Element Sets. That's all there is to it. There are just a couple more things to note:

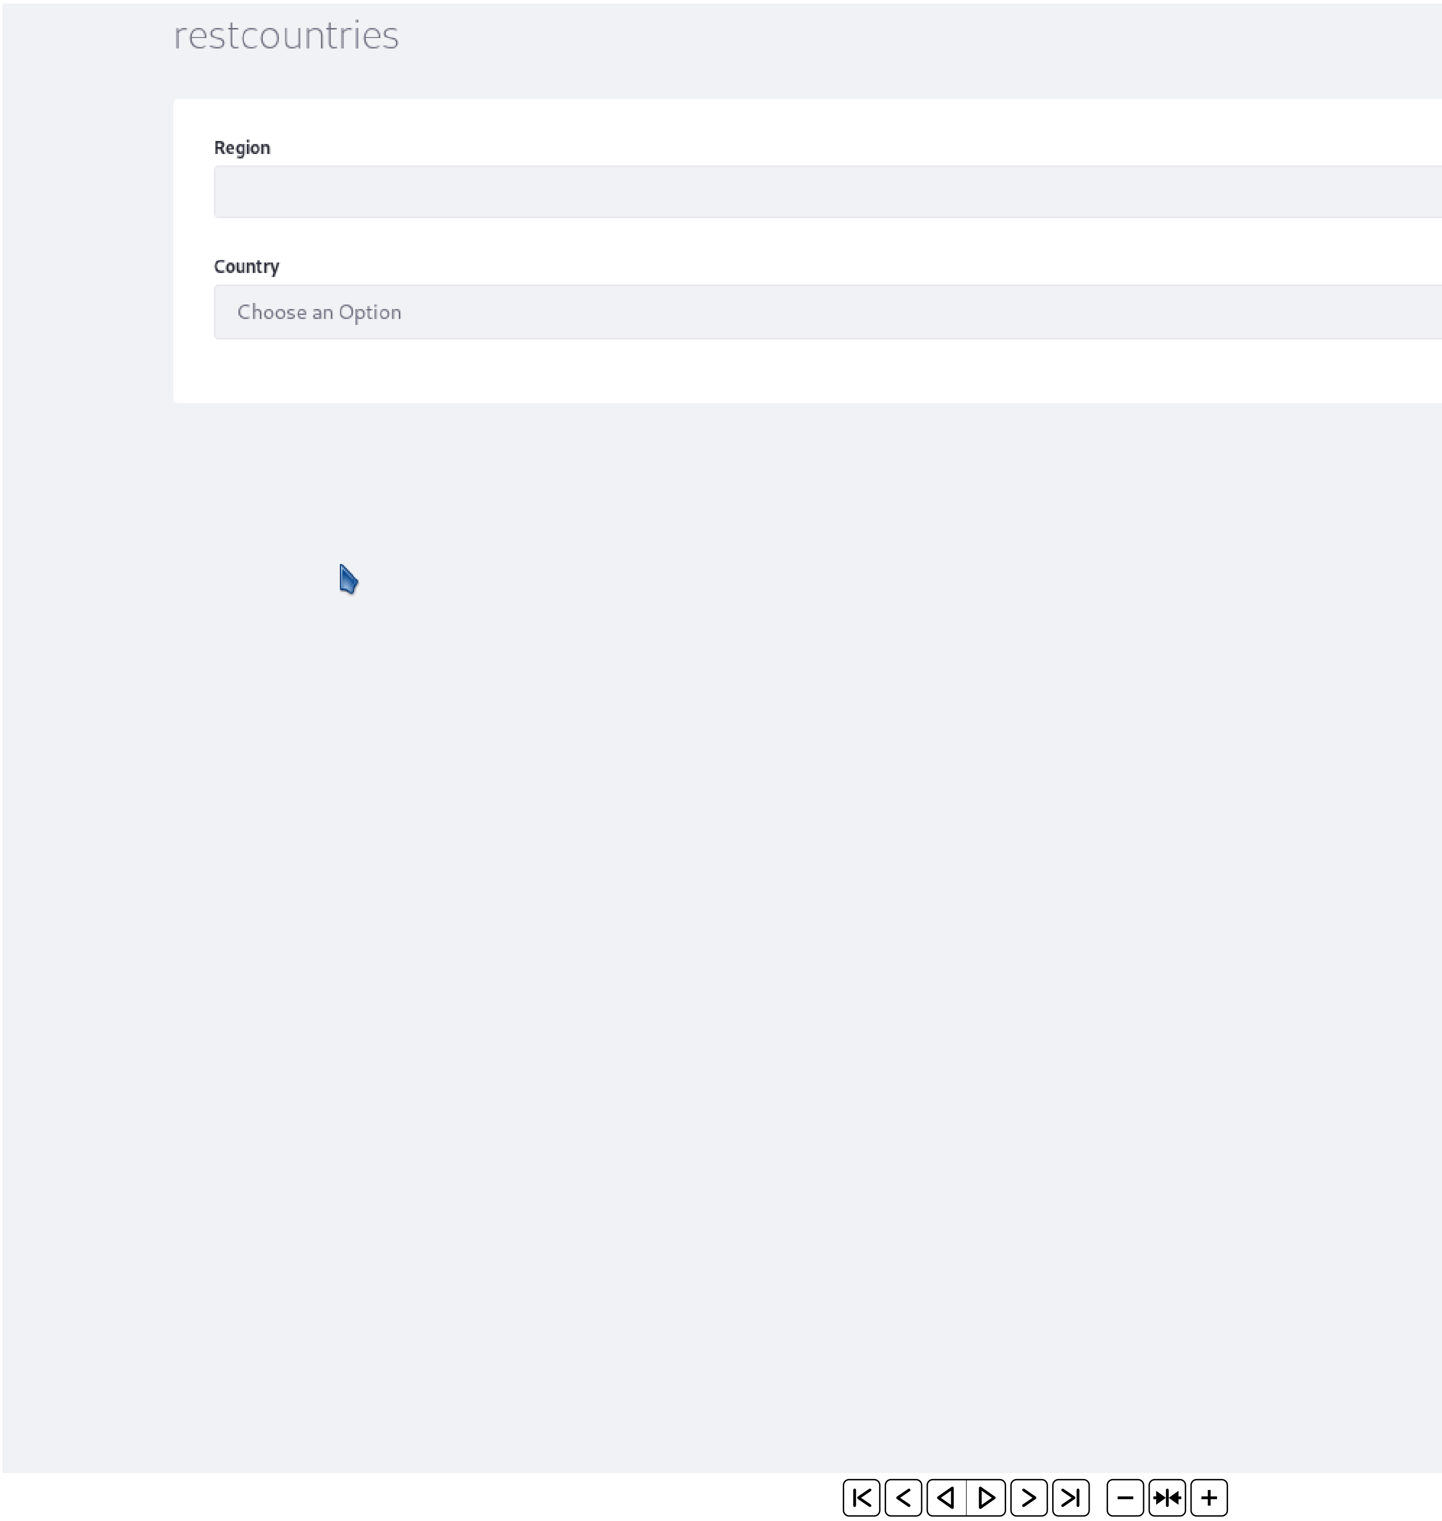

Figure 60.15: Filter countries by region of the world.

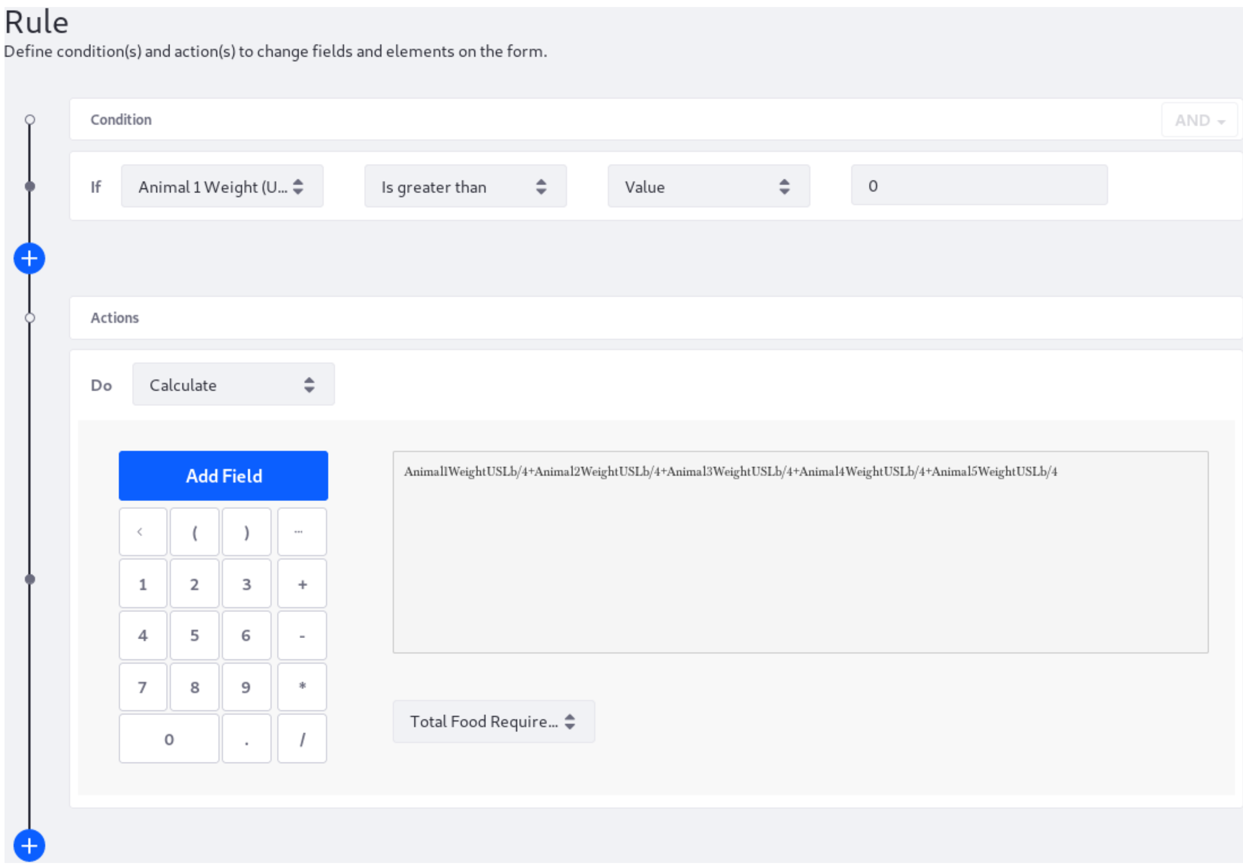

Figure 60.16: Build calculate actions with a handy calculator.

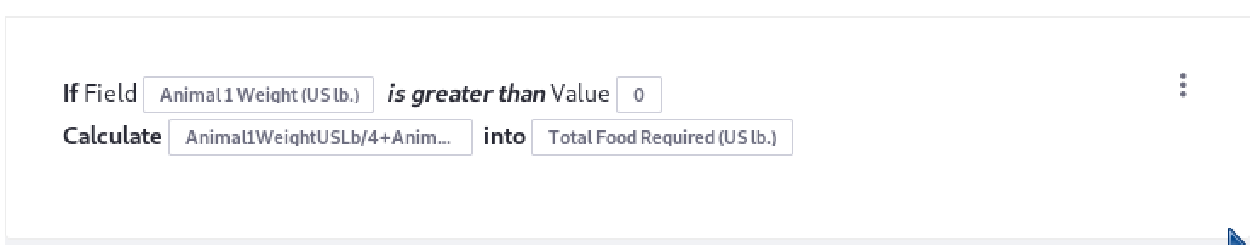

Figure 60.17: Once a rule is saved, it is displayed so that you can easily understand what it does.

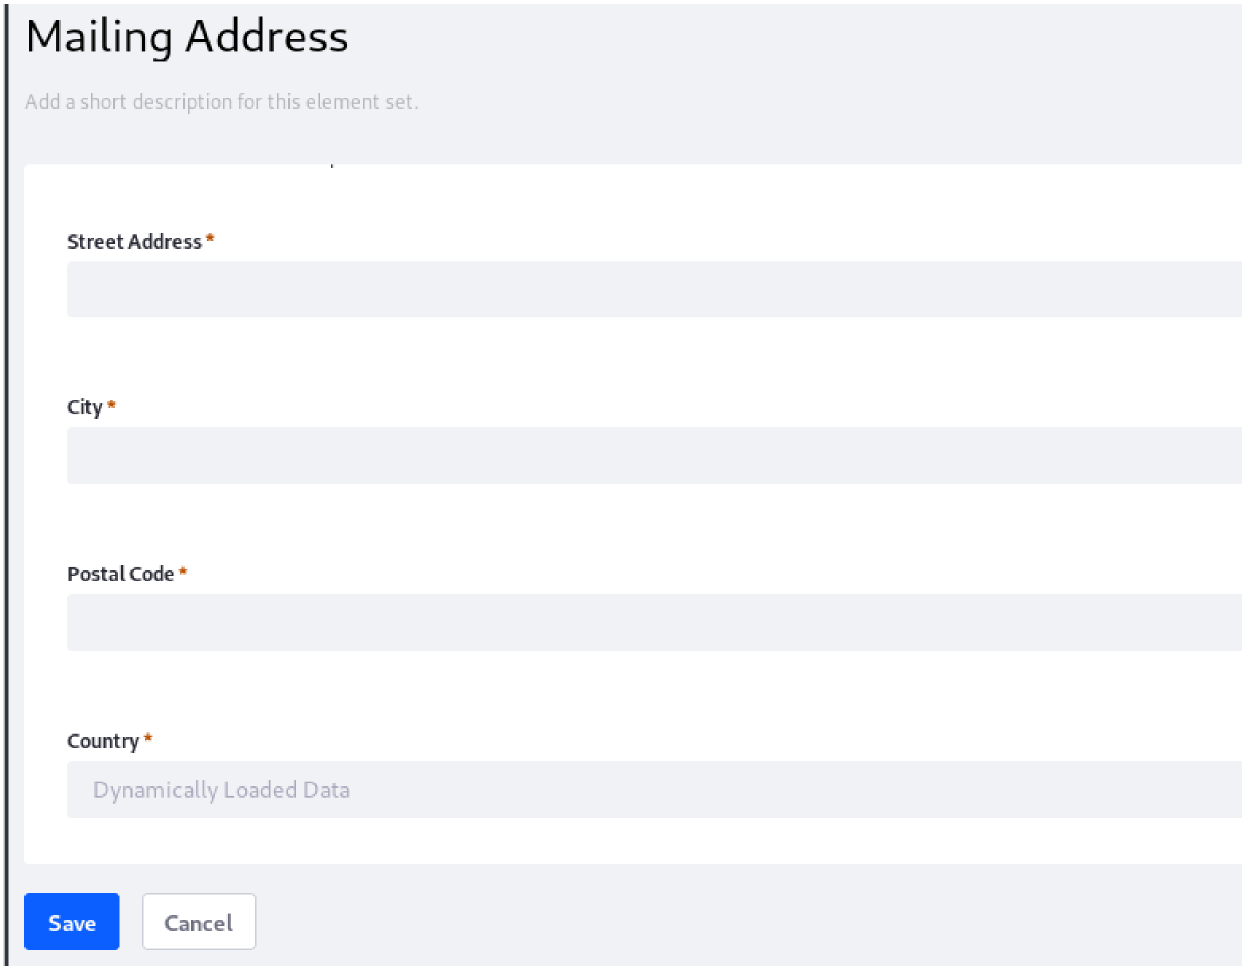

Figure 60.18: Creating Element Sets is just like creating Forms. You just can't publish them.

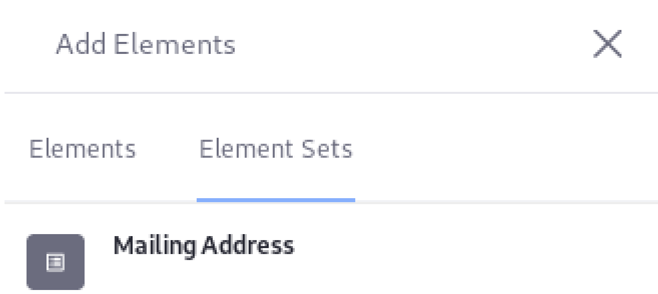

Figure 60.19: Add an Element Set the same way you add other Form Elements, like fields.

- Once an Element Set is added to a Form, there's no connection with the root Element Set. You're free to move or configure the Fields and Elements however you want.
- Editing an Element Set doesn't retroactively affect the Forms where the Element Set was used.

Think ahead. Are there some common fields you'll commonly need to configure in your Forms? If so, create them as Element Sets once and save yourself repetitive work.

60.11 Data Providers

Select from List fields can hold a lot of options. There are around 200 countries on Earth, for example. If you have unoccupied unpaid interns you could ask them to type each country into the Select from List field manually, or you could auto-populate your select fields using a REST web service. This saves you (or your interns) the trouble of typing all those options, and you can rely on someone else (hopefully a trustworthy expert) to keep the data updated.

When setting up a data provider, you're accessing a REST web service. Use the JSON web services registered in Liferay, or any other REST web service you can access. To find a list of the ready-to-use registered JSON web services in Liferay DXP, navigate to http://localhost:8080/api/jsonws (assuming you're running a local server). Browse the available Liferay services. Many times, the services useful to you in the Forms application get a list of something. Find the `get-countries` JSON web service (there are two—use either one) and click on it, then click *Invoke*. The *Result* tab shows a list of countries using JSON syntax, like this:

```
[
  {
    "a2": "AF",
    "a3": "AFG",
    "countryId": "20",
    "idd": "093",
    "mvccVersion": "0",
    "name": "afghanistan",
    "nameCurrentValue": "Afghanistan",
    "number": "4"
  },
    ...
```

That's the record for the country Afghanistan. As you can see in the *URL* tab, the URL you entered into the data provider form is the same as the one generated for accessing the `get-countries` JSON web service. Find the URL for any registered JSON web service using this same procedure.

Note the field you want Users to select. With this service, it's most likely `nameCurrentValue`, because it contains the full, properly capitalized name of the country.

Adding a Basic Data Provider

To add a *Countries of the World* Data Provider for use in your Forms,

1. Go to the Forms application.
2. Click the Options button (⋮) and click *Data Providers*.
3. Click the Add button (+).

   The REST Data Provider form loads.

4. Fill out the Name and Description fields.

   Name: `Countries of the World`

5. Enter the URL and authentication tokens for the REST service. For the `get-countries` service:

   URL: `http://localhost:8080/api/jsonws/country/get-countries/`

   User Name: `test@example.com`

   Password:`test`

6. In the Outputs fields, specify which field from the REST service populates your select list.

   Label: `Country Name`

   Path: `$..nameCurrentValue`

   Type: `List`

7. Save the Data Provider.

What's that `$..` before `nameCurrentValue`? It's JsonPath syntax to navigate the JSON data structure and specify the path to the output. Learn more about JsonPath here and here.

Using a Data Provider in a Select Field

Once the Data Provider is configured, use it to populate a Select from List field:

1. Go to the Form Builder (add a new form or edit an existing one)

2. Drag a Select from List field onto the form.

3. In the Create List section, choose *From Data Provider*.

4. Choose the Data Provider and its Output Parameter:

   Choose a Data Provider: `Countries of the World`

   Choose an Output Parameter: `Country Name`

5. Publish the form and test it.

Your Data Provider is now being used to populate a select field. However, this form should be submitted by Guest users, who don't currently have permission to see the list of results from the data provider. Arrgh! Now what?

Granting Data Provider Permissions

To configure the data provider's permissions, go to the Forms application (*Site Administration* → *Content* → *Forms*). Open the Options menu ( ) and select *Data Providers*. For the data provider you want to configure, click the Actions button ( ), then *Permissions*.

Configure the permissions you need. If Guests are to fill out the form, they need the *View* permission, or else they won't be able to see the options provided by the data provider. Once you grant permissions, click *Save*.

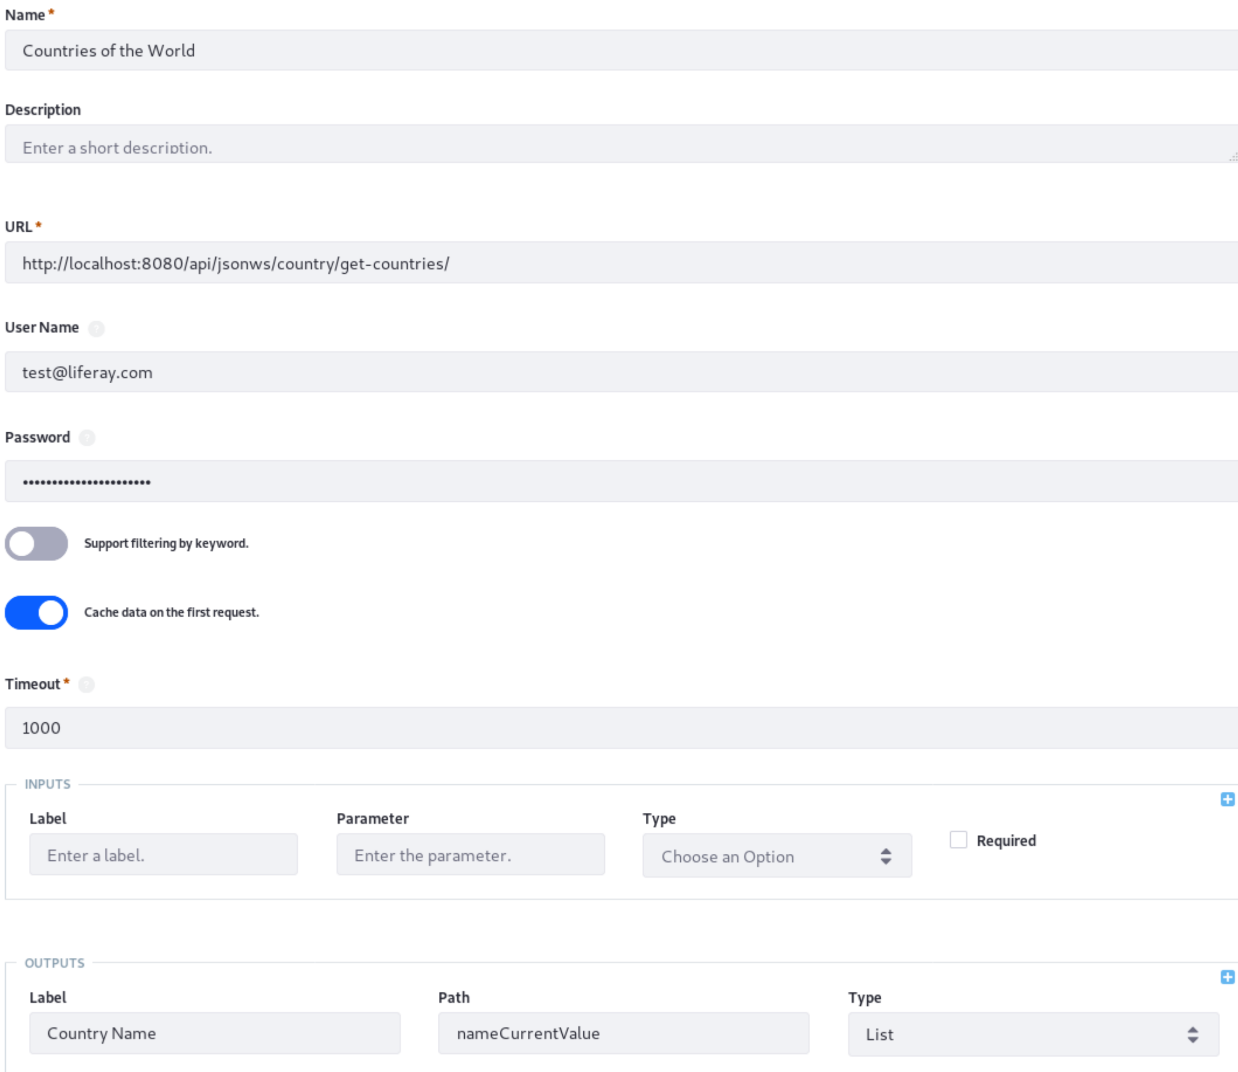

Figure 60.20: Set up a simple data provider in no time.

Data Provider Configuration

The above instructions cover adding a basic Data Provider. Knowing more about each field in the Data Provider setup form opens up more possibilities.

URL The URL of an internal or external REST service endpoint. Consider the REST service at https://restcountries.eu/, which contains a REST API endpoint to find countries by `region`:

```
https://restcountries.eu/rest/v2/region/{region}
```

Data Provider URLs can take two parameter types: path parameters and query parameters.

Path parameters are part of the URL calling the REST web service, and are added using the pattern `https://service-url.com/service/{path_parameter_name}`:

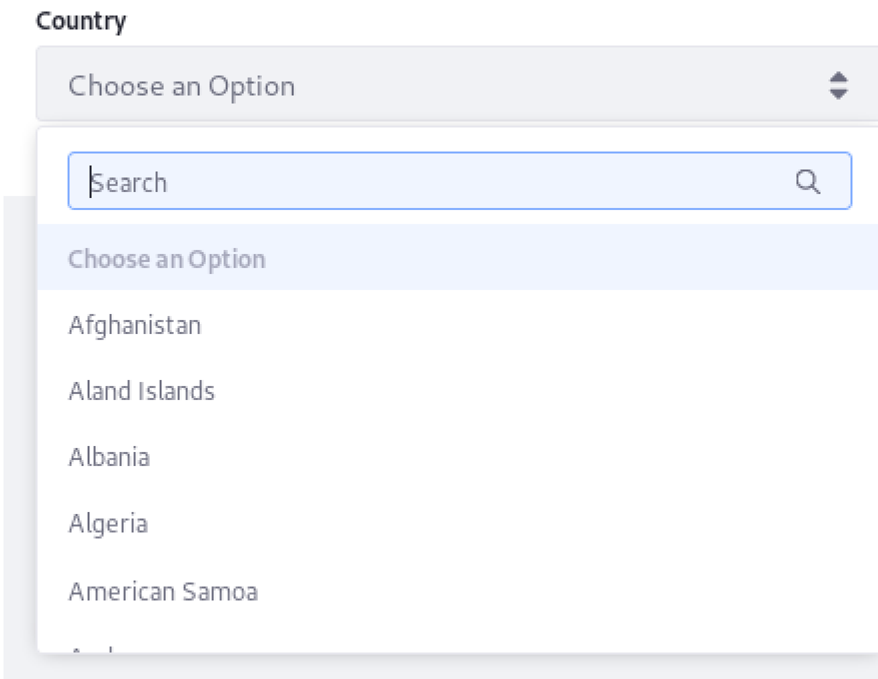

Figure 60.21: Form users select an option form the list populated by the Data Provider.

The `restcountries.eu` service's `region` endpoint's path parameter is `{region}`. Path parameters are mandatory parts of the URL, so make sure you specify an Input (see below) with a *Parameter* field value matching the path parameter from the URL.

Query parameters are complementary parts of the URL that filter the output of the service call, following the pattern `?query_parameter=query_parameter_value`:

```
https://restcountries.eu/rest/v2/all?fields=capital
```

Unlike path parameters, query parameters are optional.

User Name and Password Credentials used to authenticate to the REST Web Service, if necessary.

Cache data on the first request. If the data is cached, a second load of the select list field is much faster, since a second call to the REST service provider is unnecessary.

Timeout The time (in ms) to allow the REST service call to process before aborting the request, if a response is not returned.

Inputs Configure path or query parameters from the REST service to filter the REST service's response. Specify the Label, Parameter, and Type (Text or Number), and choose whether the input is required to use the Data Provider. You can add multiple Inputs. To provide a way for users to specify the input value, use an *Autofill* Form Rule. A User enters input into one field, and their input is sent to the REST service. The REST service's response data is filtered by the input parameter.

Outputs The Parameter to display in Select from List or Text fields with autocomplete enabled. You can add multiple Outputs. Outputs can be filtered by inputs (see above) but can also be displayed without configuring input filtering. Specify the Label, Path, and Type (Text, Number, or List). The Path field is specified in JsonPath syntax, so it must always start with a $. The type of data returned by the Path must match the type you choose in the Type field. Using the `restcountries.eu` service, specify the `name` field as an Output by entering enter `$..name` in the Path field.

If you have a more complex JsonPath expression to construct (for example, you need the names of all countries with a population over 100 million—`$..[?(@.population>100000000)].name` with the `restcountries.eu` service), consider using an online JsonPath evaluator, like this one or this one.

Hint: To display one value to the user, but persist another in the database, enter both into the Paths field, separated by a semicolon:

```
`$..name;$..numericCode`
```

If this is used with the `restcountries.eu` data provider, the name of the country is displayed to the User, while the numeric country code is stored in the database.

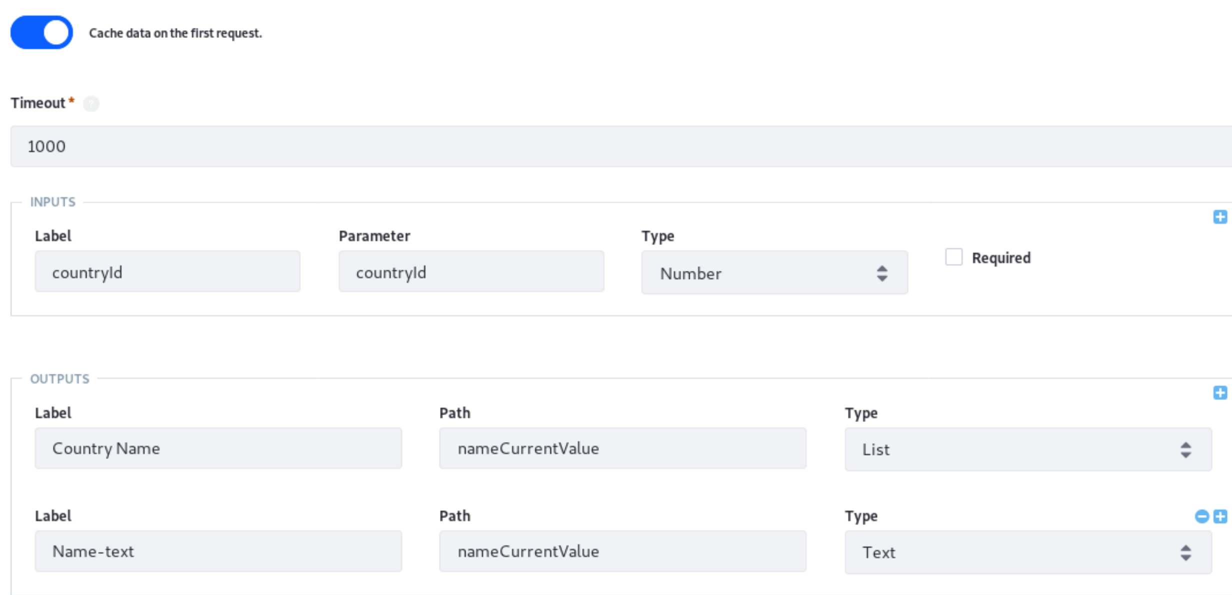

Figure 60.22: Set up Data Providers to display data retrieved from a REST service.

Troubleshooting Data Provider Errors

To uncover errors arising from Data Provider failures, configure log levels for these services:

Category: `com.liferay.dynamic.data.mapping.data.provider.internal.DDMDataProviderInvokerImpl` *Level:* WARN

Category: `com.liferay.dynamic.data.mapping.form.field.type.internal.DDMFormFieldOptionsFactoryImpl` *Level:* DEBUG

With Data Providers, the world's (RESTful) data is at your disposal to use with the Forms application

60.12 Auto-Save

Losing progress on a partially created form is bad. Make sure to save your work frequently as you're creating forms. But if you forget to save your work, Liferay Forms has your back.

By default, a form is auto-saved every minute. You won't notice anything in the form builder while the back-end auto-saves the form. You can change the auto-save duration in *Control Panel* → *Configuration* → *System Settings* → *Forms* (in the Content section). To disable auto-save, set the interval to *0*.

For unpublished forms, an auto-save works just like a manual save. For a published form, however, auto-saved data isn't automatically propagated to the form. You must click the *Save* button in the form builder to publish the changes.

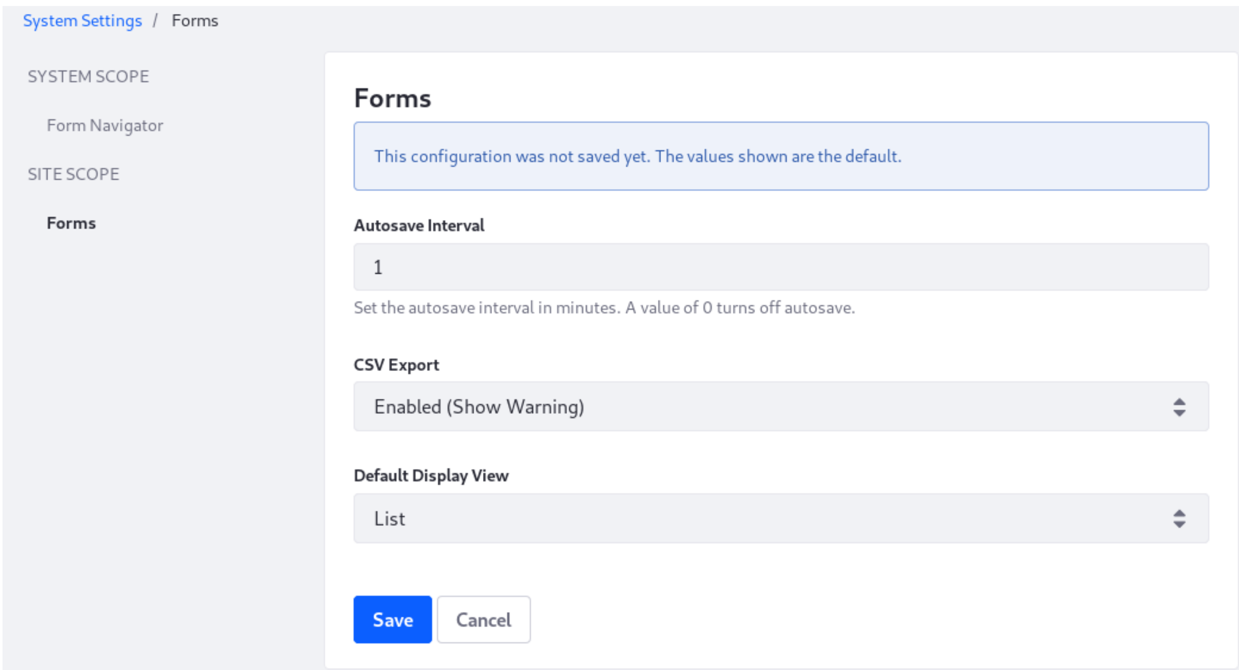

Figure 60.23: Configure the auto-save duration.

60.13 Translating Forms

Forms can be translated to any locale that Liferay DXP supports. The form builder specifies a translation of the form's default language.

The form's default language and the available translations are set in the site's language configuration.

Follow these steps to create a form translation:

1. Go to *Site Administration* (your site's menu) → *Content* → *Forms* and open the form to translate.

2. Click *Add Translation* and choose from the available languages.

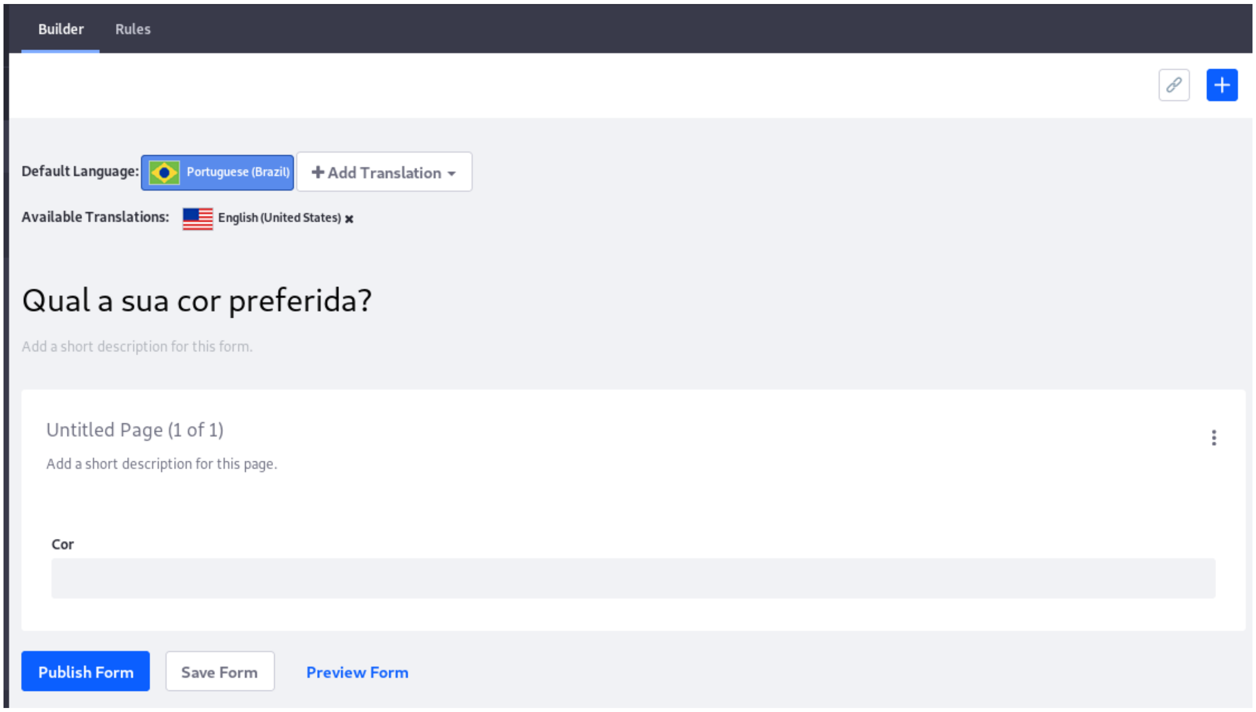

Figure 60.24: A form is translate-able into any supported language.

3. Translate the form's title, field labels, field options, field placeholder text, and any other text visible to the user.

4. Save and publish the form.

To fill out a translated form in a translated language,

1. Access the form. If a signed-in user accesses the form and a translation is available in the user's language, the user sees the translated form by default.

2. To see the form in a different language, click the language icon and select a language.

3. Fill out the form as usual and click *Submit*.

Note: Translations work differently depending on how a User accesses a Form:

1. If accessed in the Form widget on a Liferay DXP page, the Form is displayed in the User's language automatically. If there's no translation available for the User's language, the default language of the Form is displayed.

2. If accessed via direct URL, the Form translation must be selected manually.

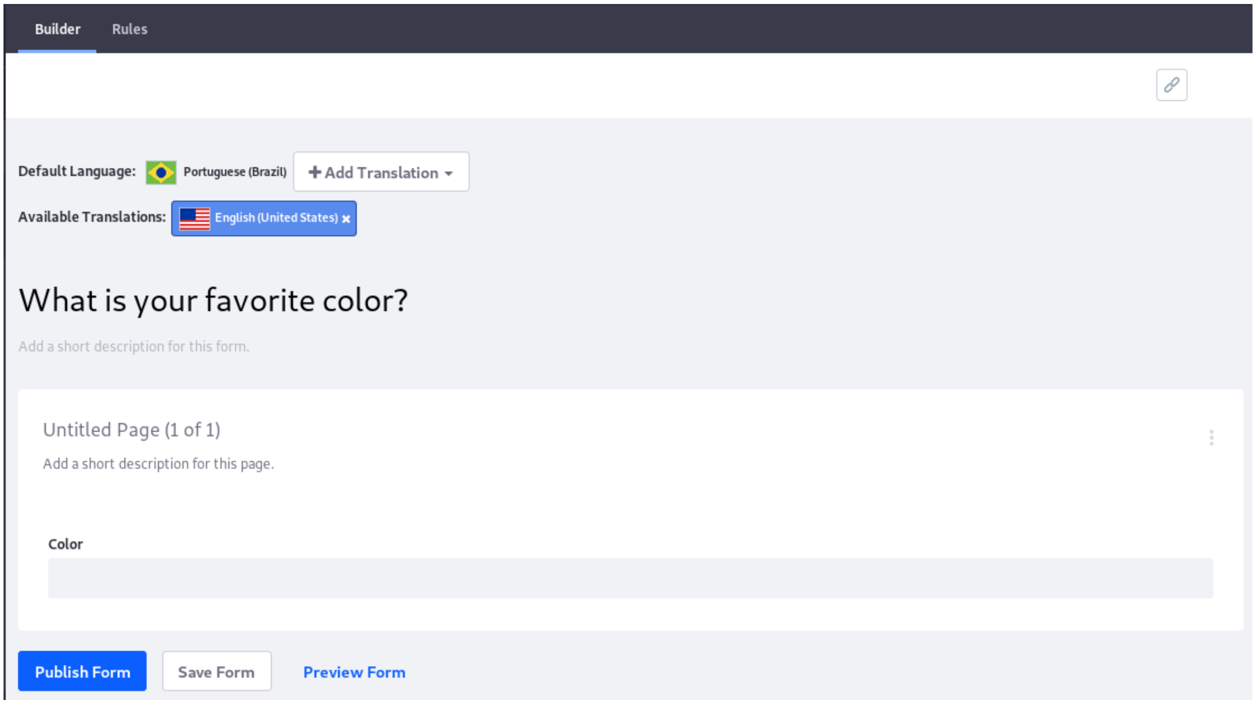

Figure 60.25: Translate as much of the form as possible into each language you expect users to need.

Figure 60.26: Select the form's language.

60.14 Autocompleting Text Fields

It's been scientifically proven that Internet users are lazy (not you, of course—other Internet users). For example, some users may not fill out your form if you make them type the entire title of their country in an employment application. This is especially true if they're filling out the form on their mobile devices. In this article, you'll learn to make users' lives easier by configuring autocomplete on a form's text fields.

Why not just use a select field with a data provider to guide user input? Sometimes a data provider can't encompass all possible field entries. For example, if your data provider doesn't include mythical countries founded on old sea platforms, users from that country wouldn't be able to enter anything into the select field. Alternatively, you can use a text field with autocomplete so users can begin typing their country's name and then select it from a list when it appears. Autocomplete combines a text field (accepting any response that meets your validation criteria) and common choices to select from. It's a win-win situation.

Configuring Autocomplete

Before configuring autocomplete for your text fields, create a form and add a text field. If you want the autocomplete options to be populated by a REST data provider, configure one before creating your form. Now you're ready to configure autocomplete for the field:

1. In the field configuration sidebar, click the *Properties* tab.

2. Click the *Autocomplete* switcher so it's enabled.

3. Select a data provider or create one manually. You can set up a data provider from a REST service, or manually enter the options users should see when they start typing in the text field.

4. Save and Publish the form.

Once users begin entering text into the field, a selection list of options appears. As they enter additional text, the list is refined to include only options that contain the currently entered text. For example, the imaginary users from Sealand (all two of them) begin reluctantly typing their country of origin by entering an *S*. They're delighted to see a selection list with a bunch of countries containing the letter *S* appear for their selection convenience. If they continue typing and enter *e*, the list is refined to options that have *se* in their name (for example Serbia and Senegal). If they continue typing and enter *a*, they'll now only see one option, Sealand, if it's in your data provider. Selecting it from the list after typing the first three letters is much easier than typing the remaining letters.

What will the Forms team think of next? Configuring telepathic connections to the Forms application would be nice. Then users could just think their form field entries into existence. Stay tuned.

60.15 Form Success Pages

After users submit one of your whiz-bang forms, what's next? How will they know they're done and can close the browser window or tab? What if they think their submission didn't go through

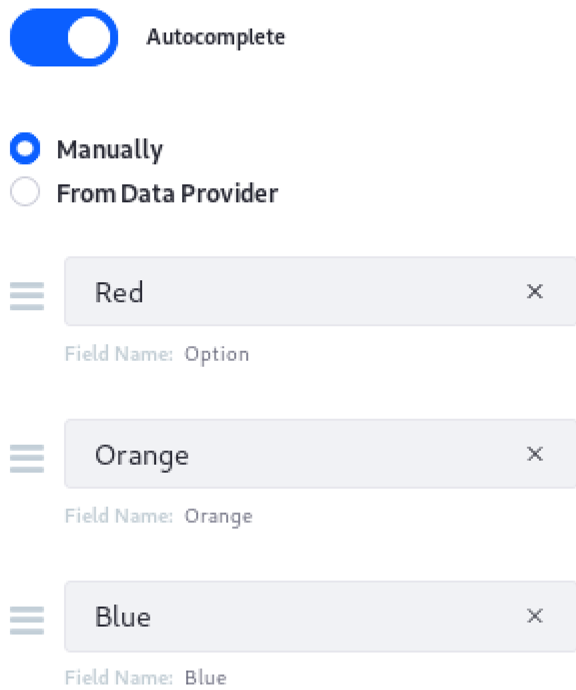

Figure 60.27: You can configure a manual data provider to specify the options users can select from.

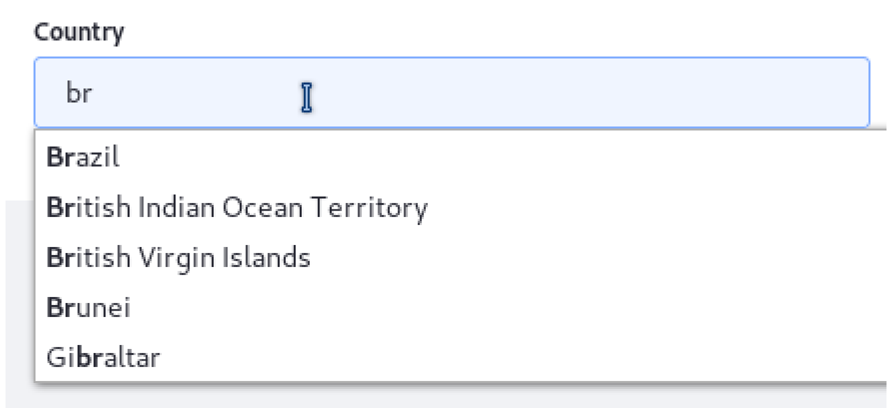

Figure 60.28: When typing in a field with autocomplete, users are presented a list of selections from the configured data provider. The displayed results are filtered to include only selections containing the text entered by the user.

and wonder if they need to fill out the whole form again? By default, submitting a form displays the default success message and returns users to the form's now empty first page. Don't leave users feeling equally empty. Instead, configure a *Success Page*. A Success Page is a terminal page showing users they've finished filling out the form and their submission was successfully received. A Success Page can even urge users to close the browser window or tab.

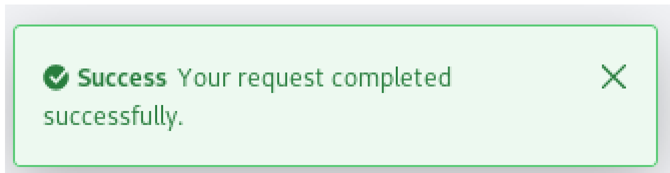

Figure 60.29: The default success message alerts users when their request completes successfully.

A Success Page is simple. It has a title in bold text and a description beneath the title. A common alternative to using a Success Page is to redirect users to a different page in your Site. What should you put in a Success Page? Whatever you want. If you can't think of anything important or creative to say, use the default message:

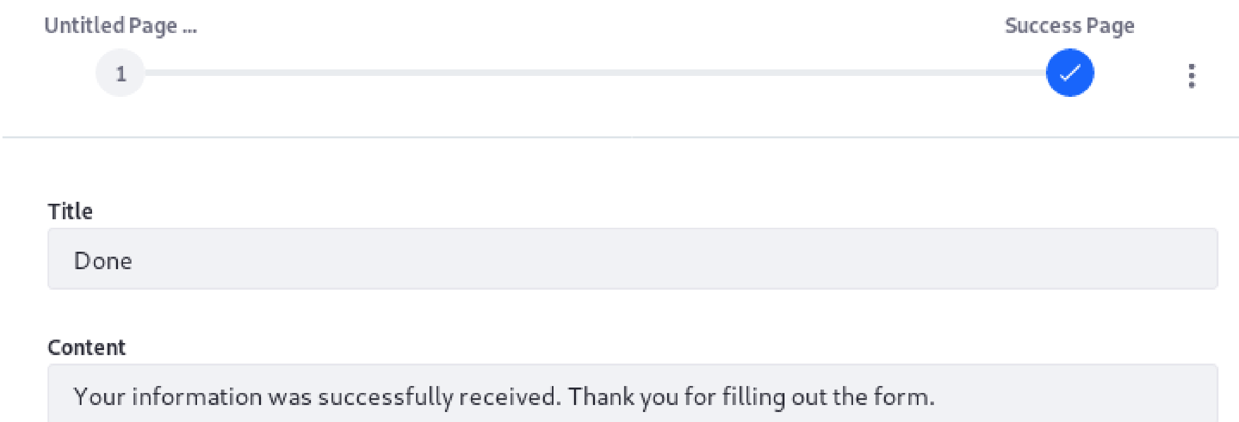

Figure 60.30: There's a default Success Page message if you can't think of anything else to say.

To configure a Success Page,

1. Add a form in *Site Administration* (your site's menu) → *Content* → *Forms*.
2. Click on the form page's *Actions* button (⋮) and choose *Add Success Page*.

Once the Success Page is added to your form, fill in the *Title* and *Content* fields however you please. When the form is saved and published, the Success Page is live for your form users.

Note: You can't preview the Success Page. Success Pages can only be viewed once a form is submitted, and the *Submit* button isn't available in the form preview. The *Preview Form* link in the form builder only lets you preview the form's regular pages (use the *Next* button to navigate through the form).

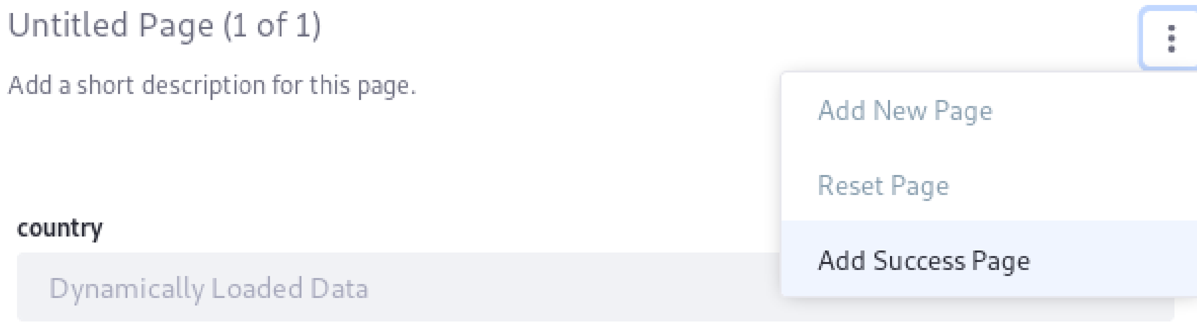

Figure 60.31: Add a Success Page using the edit menu for the form page.

To see what your Success Page looks like, submit a test entry of the form and then delete it if needed. For more information, see the documentation on viewing and managing form entries.

60.16 Workflow and Forms

Kaleo is a workflow engine for sending a submitted asset through a workflow process before it's published. Most assets are configured to use workflow at the instance or Site level.

Forms is different, and that's why its assets don't appear in the above image. There are so many use cases for forms, and there could be so many per site, that a site- or instance-scoped workflow configuration won't serve your needs well. Instead, configure workflow for *each form* separately.

Enabling Workflow in a Form

To enable workflow in a form,

1. Open the form's editor by opening the Menu, selecting your Site, navigating to *Content* → *Forms*, and clicking on the form you want.

2. Click the Options button (⋮) and choose *Settings*.

3. The Settings window has a *Select a Workflow* drop-down. Find the workflow you want, select it, and then click *Done*.

Testing the Workflow

Test the workflow process:

1. Add the form to a page.

2. Click *Submit for Publication* to submit the form entry.

Next go find the form entry in the Forms widget:

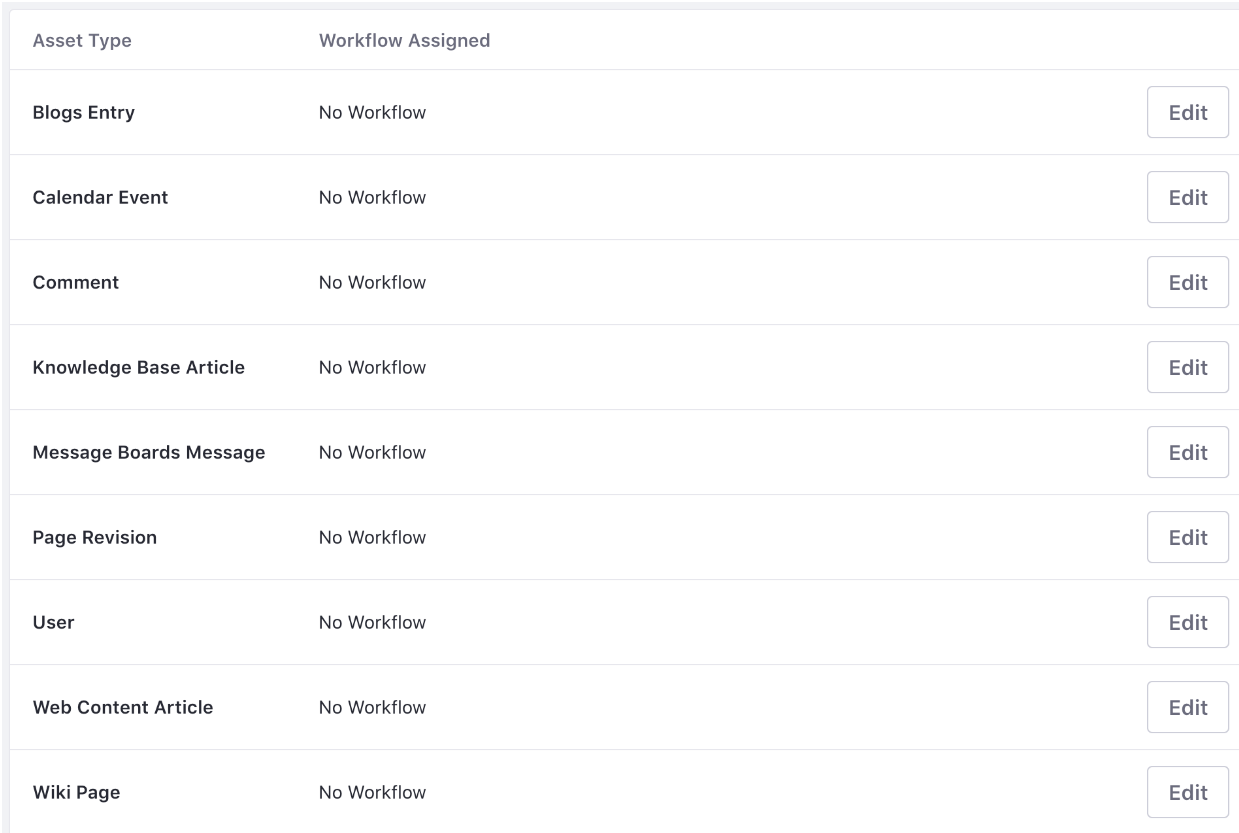

Figure 60.32: Workflow is enabled in the Control Panel or in Site Administration for most Liferay DXP assets.

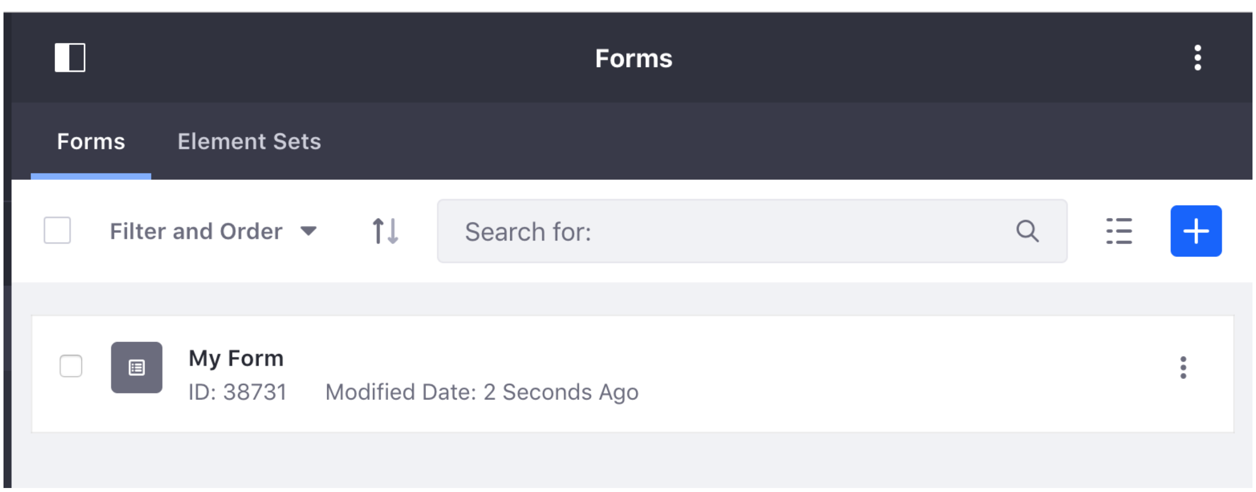

Figure 60.33: Navigate directly to a form to enable workflow.

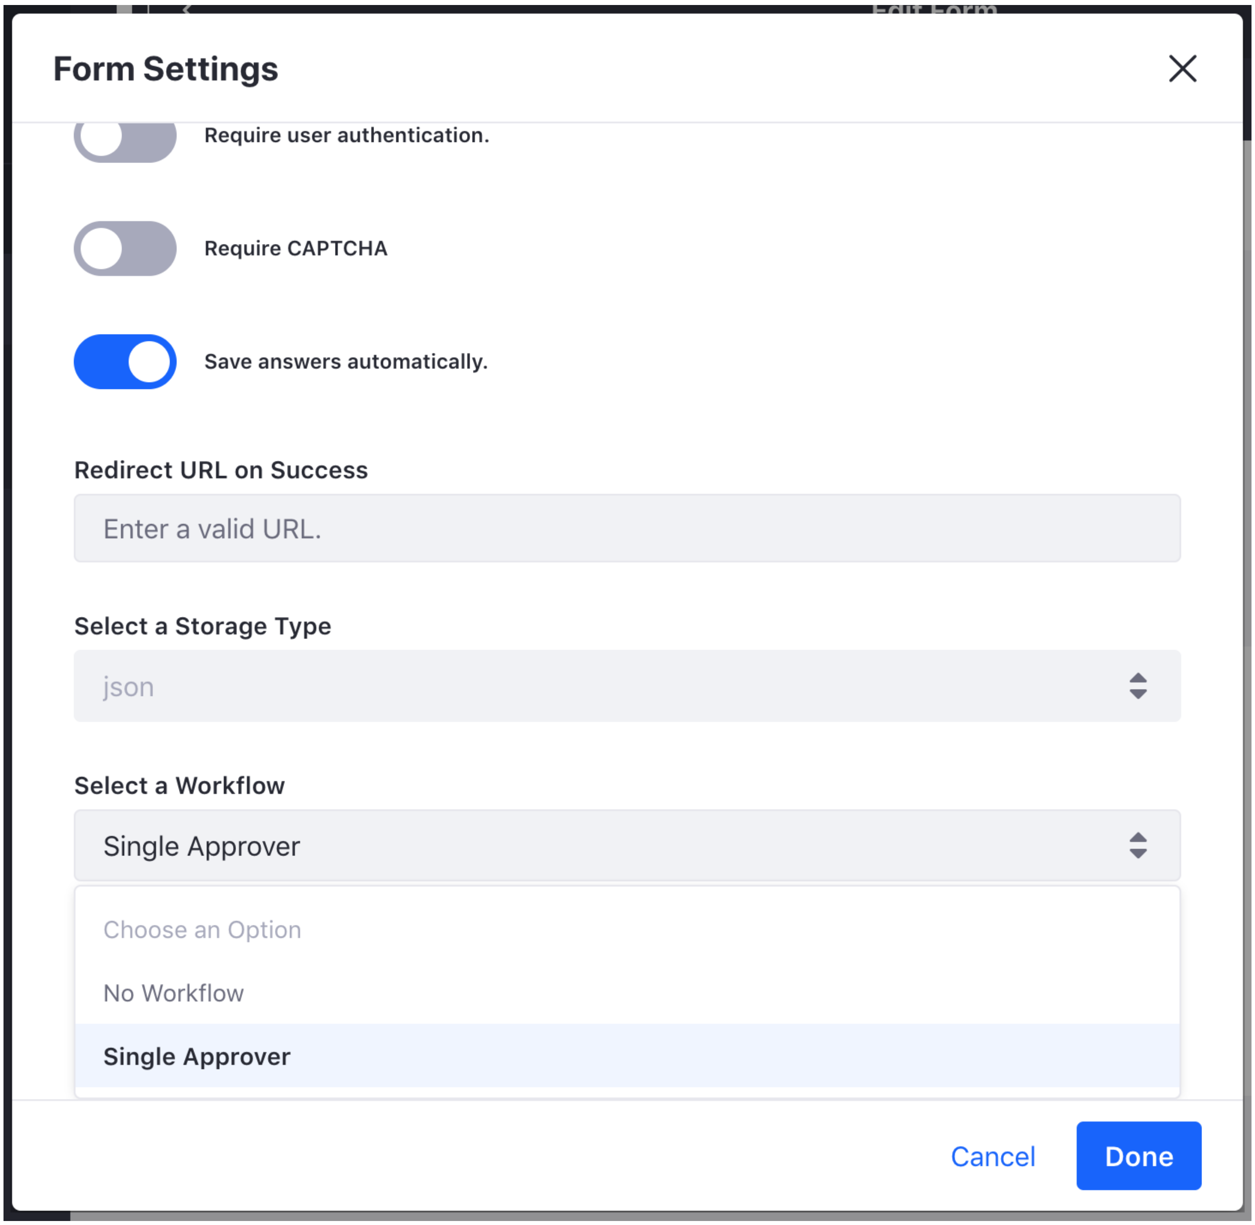

Figure 60.34: Enable workflow for each form in its Settings window.

1. Go back to the Forms application in the Menu in your Site's *Content* section.
2. Click the Form's *Actions* button (⋮) and select *View Entries*.

   The entry is currently marked *Pending*.

Now approve the form record:

1. Navigate to *My Account* → *My Workflow tasks*.
2. Click the *Assigned to My Roles* tab.
3. Click on the form entry.
4. Click the Actions button (⋮) and choose *Assign to Me*.
5. Click *Done*.
6. Click the Actions button (⋮) again, then click *Approve*.
7. Click *Done* again.
8. Navigate back to the View Entries screen for the form, and now the entry is marked as *Approved*.

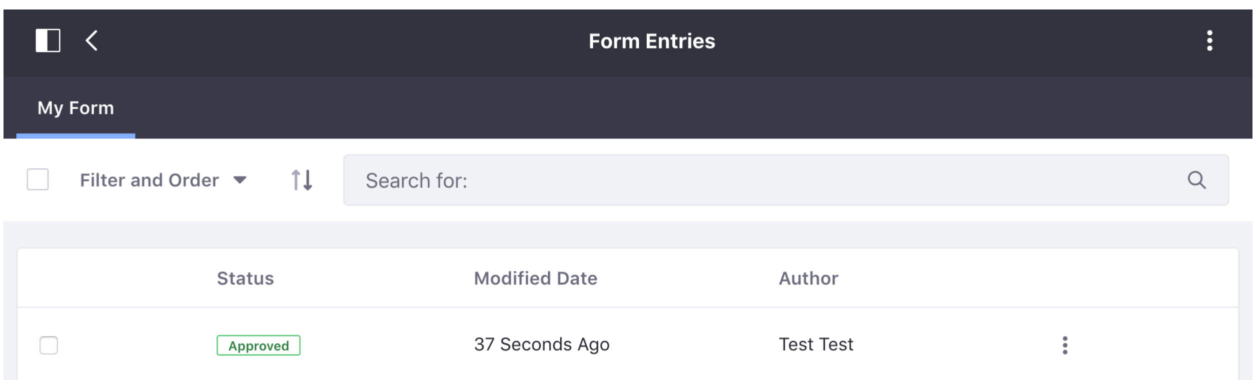

Figure 60.35: Each entry's status is visible in the Forms application's Form Entries screen.

Example Use Case: Win a Free, All-Inclusive Trip to the Lunar Resort

The Lunar Resort is giving away a free, all-inclusive trip to several lucky families. All a user must do is click the *I Want a Free Trip* link from the Site home page, to see a form they can fill out and submit in under three minutes:

The Lunar Resort wants to exclude folks who don't meet certain qualifications:

- Exclude families with more than five persons in the household.
- Prioritize families who've never visited The Lunar resort.
- Exclude anyone who's on the resort's naughty list, appropriately titled "The Dark Side of the Moon."

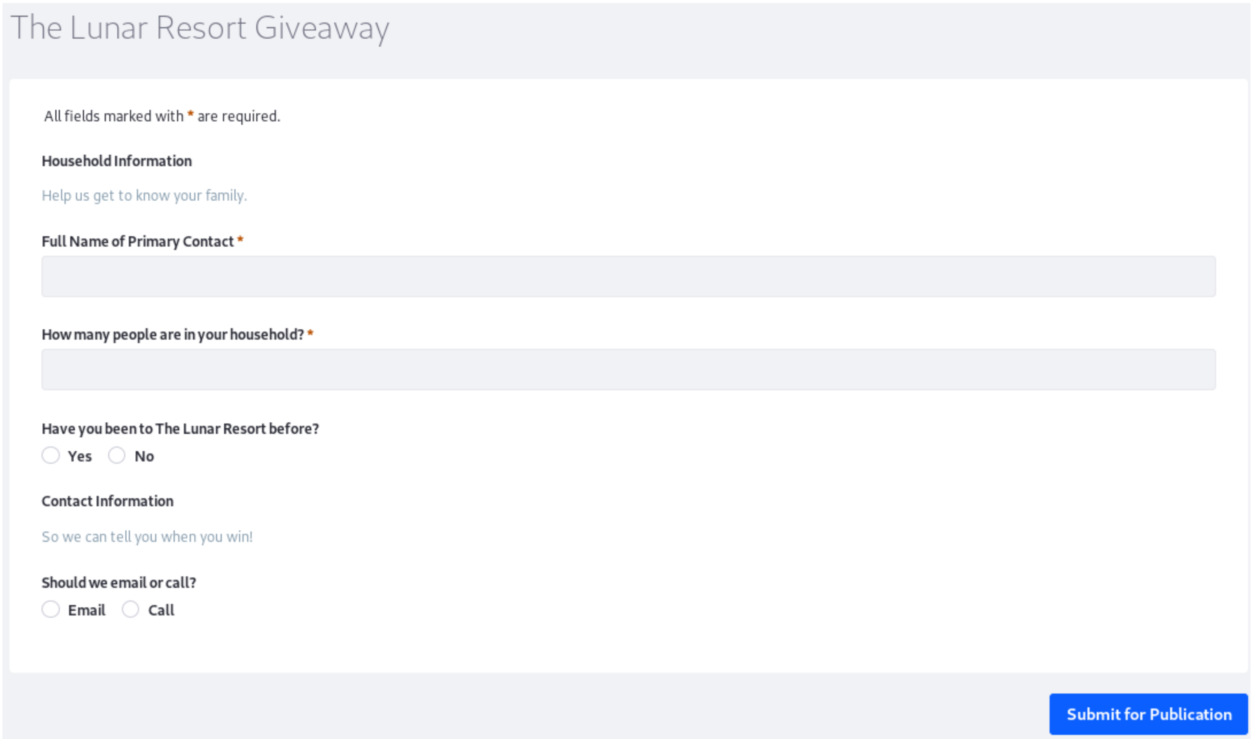

Figure 60.36: The Lunar Resort Giveaway form is ready to be filled out.

To compile a list of finalists, each form entry should go through a workflow with one review task. The Single Approver workflow works nicely for this.

After creating the form in the above image, select the Single Approver workflow from the form's Settings window, as described in the above section. Publish the form, then fill out and submit an entry. The administrative user receives a notification that a form record was submitted. Go to the View Entries screen and see that the form record is currently marked as Pending. Once the review is completed, it's marked as Approved.

60.17 Duplicating Forms and Form Fields

Duplicating effort wastes time and increases the likelihood that something won't be perfectly replicated. Instead of duplicating effort, duplicate forms and form fields.

To duplicate a form,

1. Go to *Site Administration* (your Site's menu) → *Content* → *Forms*.
2. Click the *Actions* button ( ) for the form to duplicate.
3. Click *Duplicate*.

The form is duplicated and automatically named *Copy of [Original Form Name]*. Once duplicated, you can edit the form however you want. When you duplicate a form, all configurations and form rules are duplicated as well.

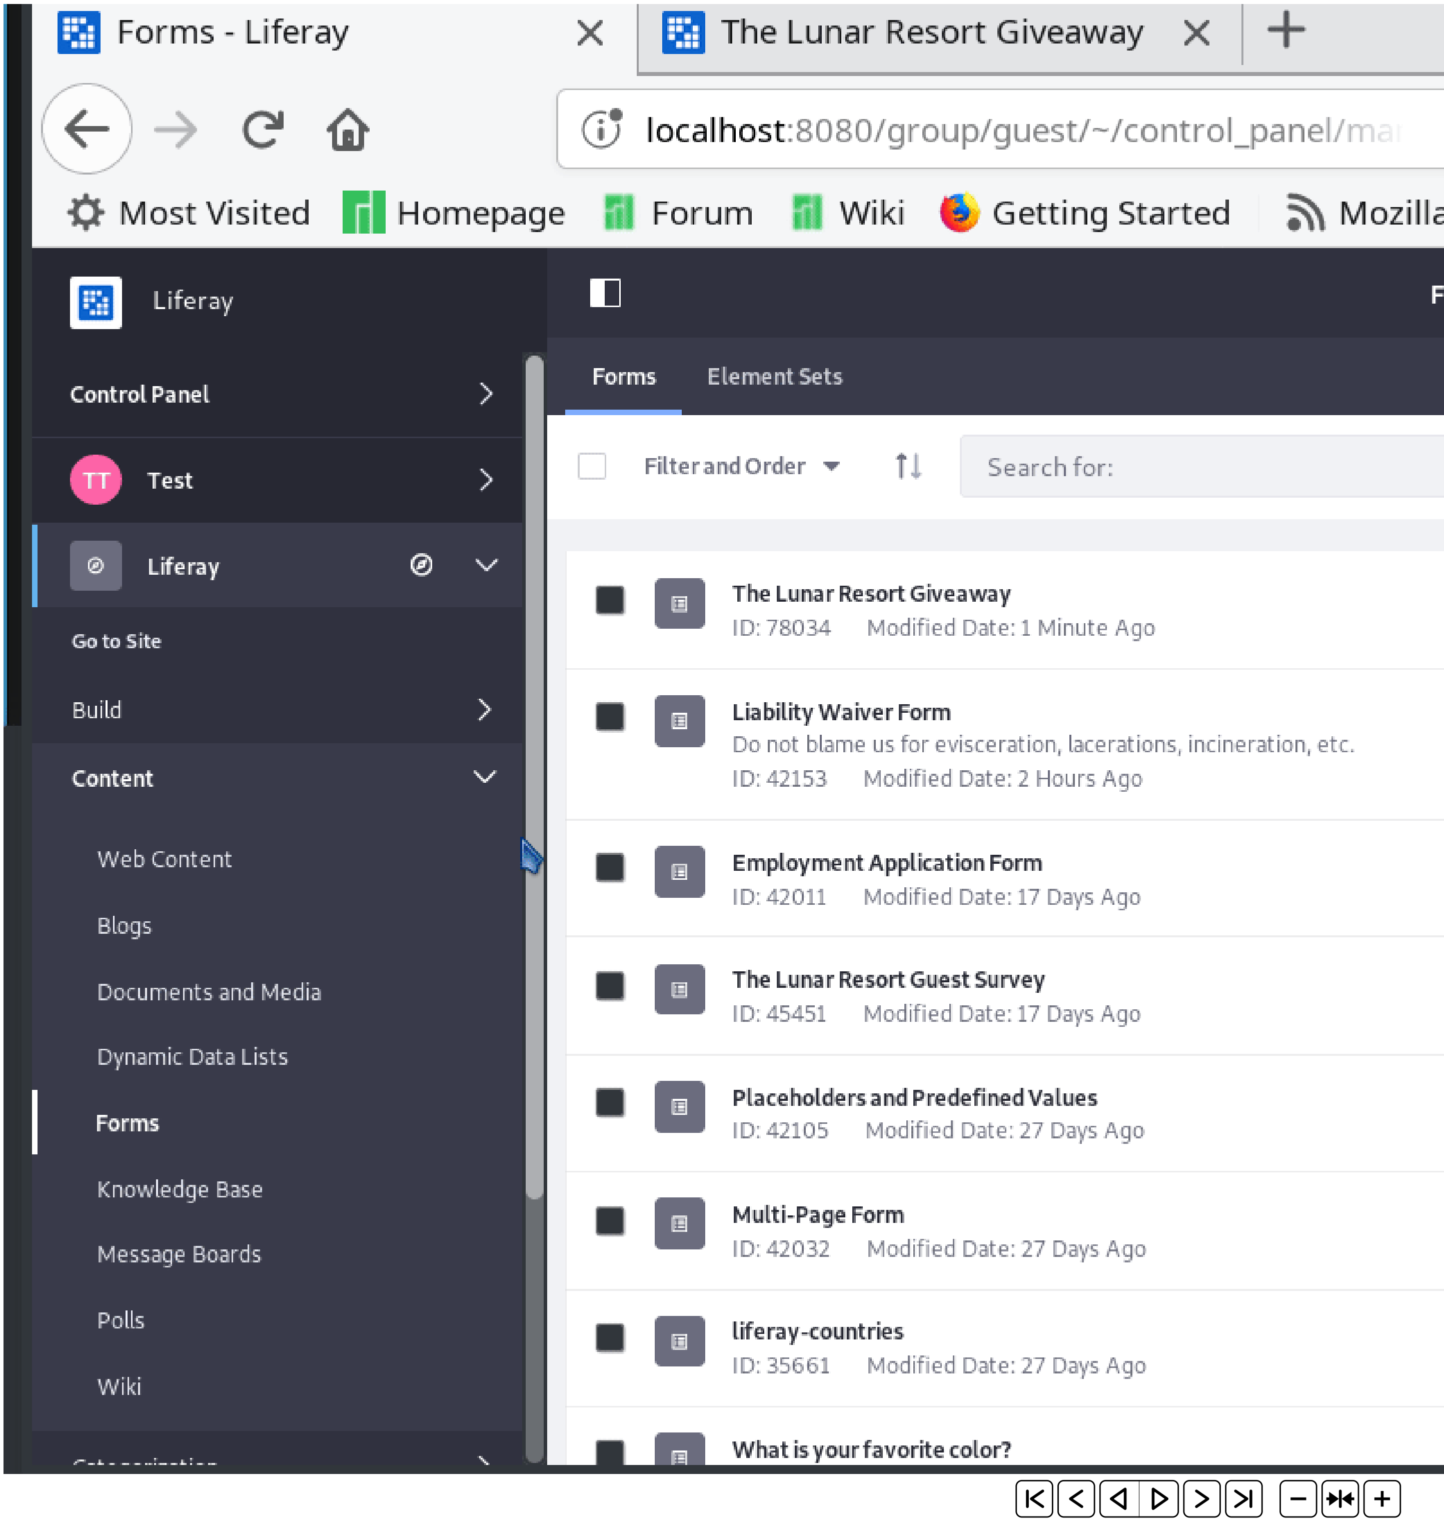

Figure 60.37: Assign a workflow to a form in several steps.

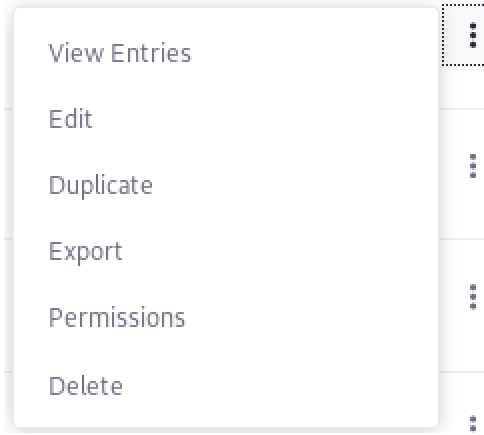

Figure 60.38: The Duplicate option works the same for forms and form fields.

To duplicate a form field,

1. Go to *Site Administration* (your Site's menu) → *Content* → *Forms* and open or add a form.
2. In the Builder view, hover over the form field to duplicate and click the *Copy* icon ( ).

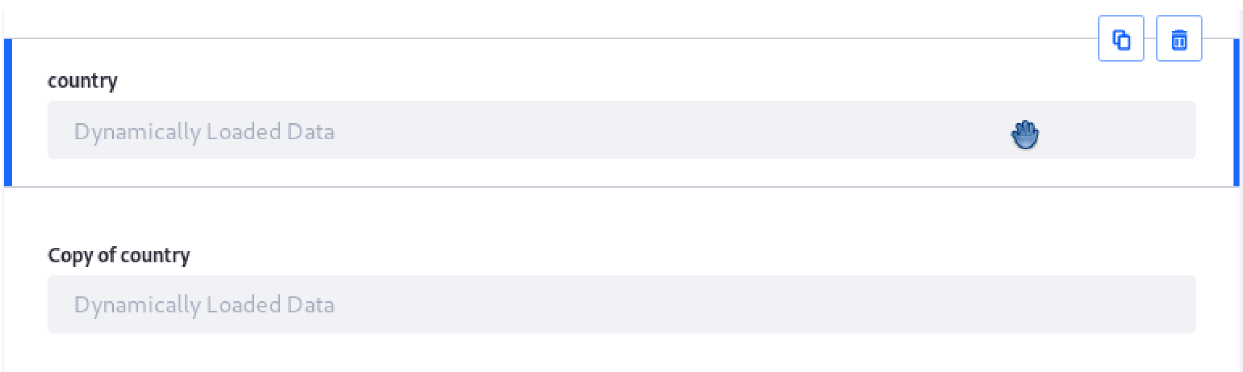

Figure 60.39: You can duplicate form fields.

The field is duplicated and labeled *Copy of [Original Field Label]*. All the form's properties, including its data provider configurations, are copied as well.

60.18 Form Pages

Are users more likely to abandon long forms with lots of scrolling? Are they more likely to see a multi-page form and abandon it without a second look, assuming that it gets longer and more tedious with every passing page? Such usability questions are worth thinking about. If you decide multiple pages are appropriate for your form, Liferay Forms supports two pagination styles: default, and alternate.

To add a form page,

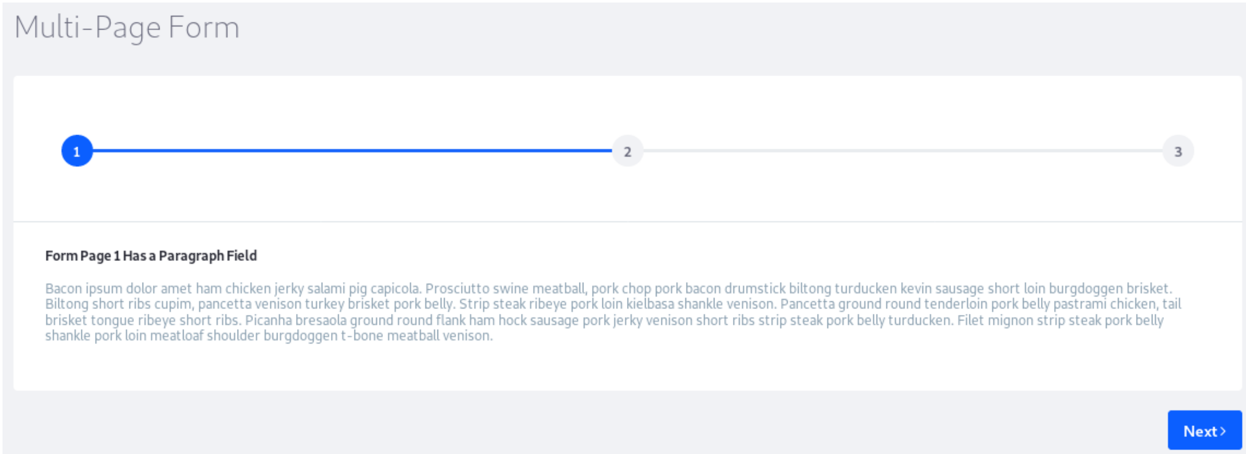

Figure 60.40: The default pagination style.

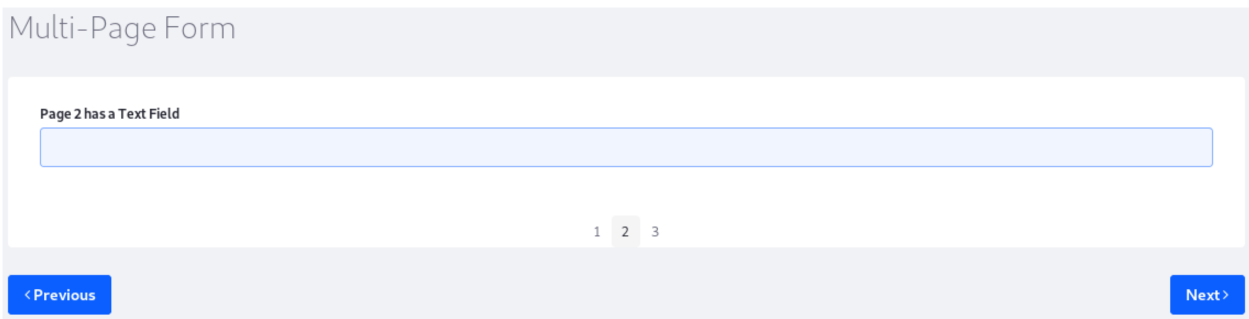

Figure 60.41: The alternate pagination style.

1. Go to a form's builder view.
2. Click the *Actions* button (⋮) at the top-right corner of the form, then click *Add New Page*.

You can also delete form pages, add success pages, and switch pagination modes.

60.19 Help Text, Placeholder Text, and Predefined Values

Form fields can have help text, placeholder text, and predefined values.

- **Help Text:** Text that appears as a sub-heading to the field label, but doesn't appear in the field entry area. Enter help text in the Basic tab of the field's sidebar menu.
- **Placeholder Text:** Text in the field entry area that isn't submitted if the field is left untouched by the user.
- **Predefined Value:** Text in the field entry area that is submitted if the field is left untouched by the user.

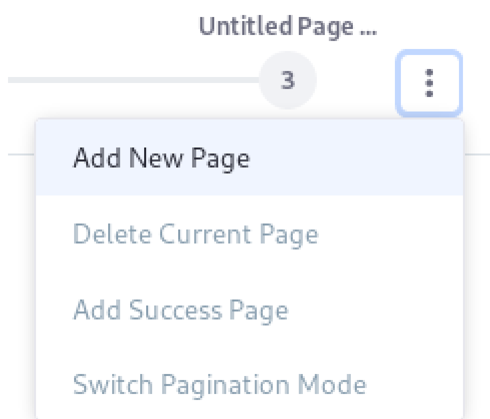

Figure 60.42: You can add new pages or reset the current page from the Page Actions menu.

All form field types can have help text, and all form field types that accept user input can have predefined values. Only text and numeric fields can have placeholder text.

To enter placeholder text or predefined values,

1. Open a field's sidebar menu.
2. Open the *Properties* tab.

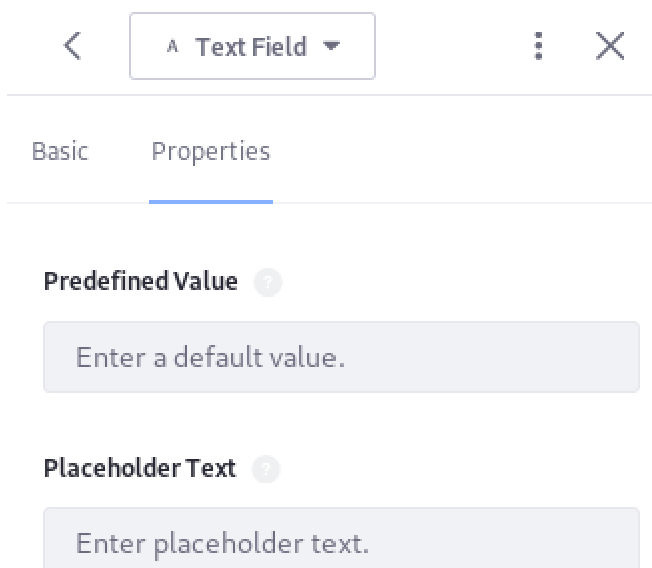

Figure 60.43: Predefined values and placeholder text are entered in the Properties tab.

For example, many forms start with a *Full Name* field. You can use help text and/or placeholder text to inform users that you need their full names, regardless of length. Alternatively, if you're asking users to specify how many sandwiches they eat for lunch, a predefined value of *1* probably makes sense.

Figure 60.44: The Full Name field here uses help text and placeholder text, while the sandwiches field uses a predefined value.

Remember, placeholder values aren't submitted if the field is left blank, so you don't have to worry about getting a bunch of submissions from *Maximillian Aurelius Piroux the 11th*.

60.20 Validating Text and Numeric Fields

Validation ensures that only certain values are entered in a field. Validation functionality is available for text and numeric fields.

To enable validation,

1. Add a Text or Numeric field to a form in the Builder view.
2. Open the field's *Properties* tab.
3. Turn on the *Validation* toggle to enable validation and open its configuration options.

Validating Text Fields

Validation for text fields contains several options. You must first choose a list of available conditions to check:

- If Input *Contains*
- If Input *Does Not Contain*
- If Input *Is not URL*
- If Input *Is not Email*

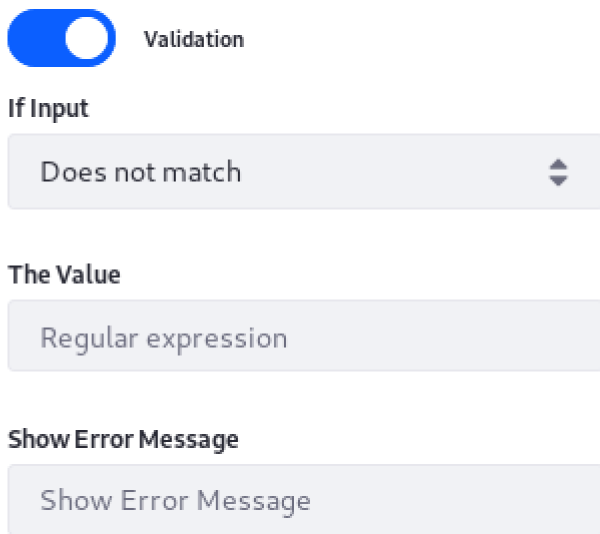

Figure 60.45: Validate data to ensure you're collecting only useful information.

- If Input *Does not Match*

If the condition isn't met, an error message is displayed to the user. Where you go from there depends on which condition you used.

If Input Contains/Does Not Contain

When you validate text data to check if it contains a certain value, there are two additional steps to take after selecting the condition:

1. Enter the text to check for.
2. Enter an error message so users understand why their submission failed.

If Input Is not URL/Email

Checking for properly formatted URLs and emails is easy. Just choose the condition from the *If Input* drop-down and enter the error message.

Valid URLs begin with `http://` or `https://`. Valid emails must contain `@`.

If Input Does Not Match

The *Does Not Match* condition is used for entering regular expressions to create custom validation criteria. For example, use this regular expression to ensure that ten consecutive numeric digits are entered in a phone number field:

```
^[0-9]{10}$
```

If you use regular expression validation, provide some explanatory text (e.g., help text, placeholder text, and a clear error message) to guide form users in entering the proper data.

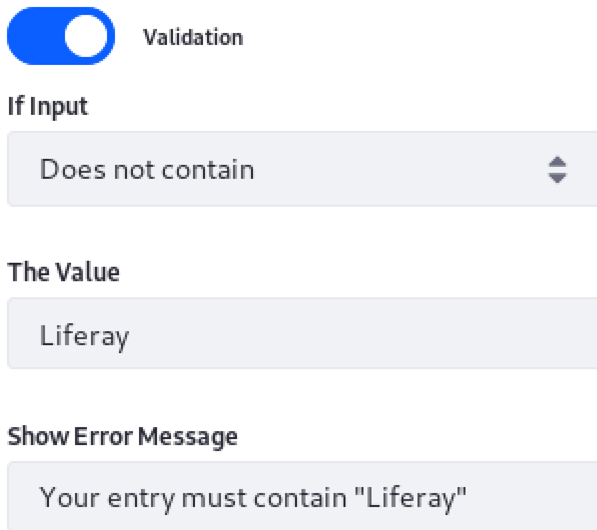

Figure 60.46: If *Liferay* isn't part of the field's value, an error message is displayed.

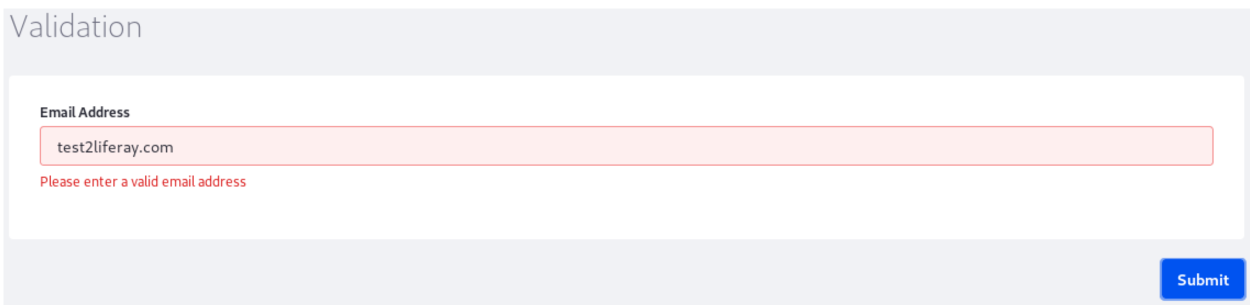

Figure 60.47: Use text field validation to make sure users enter a valid email address or URL.

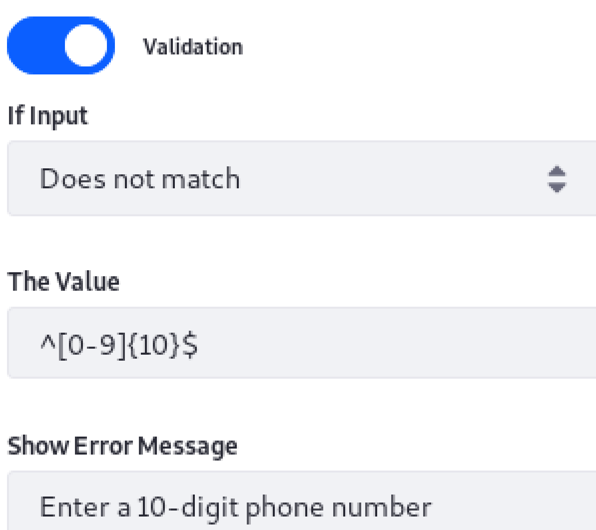

Figure 60.48: Regular expression text validation opens up countless possibilities.

Validating Numeric Fields

Numeric field validation is similar to text field validation, but the conditions compare the value of the number entered to some other value.

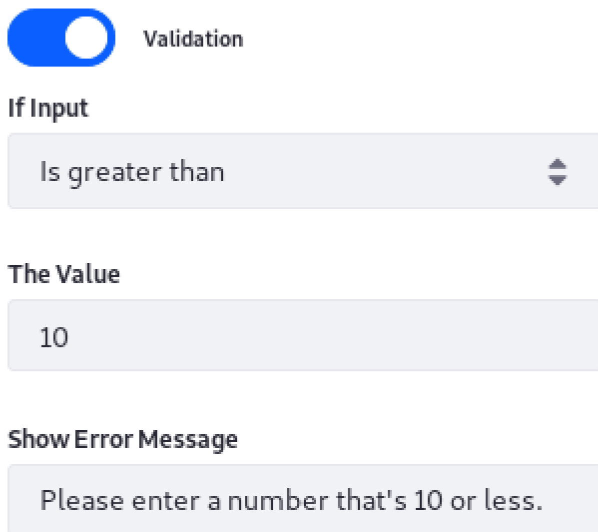

Figure 60.49: Numeric conditions constrain user-entered numeric data.

Available conditions to check include

- Is greater than or equal to
- Is greater than
- Is not equal to
- Is less than or equal to
- Is less than

For example, to make sure users don't enter a number over 10, enable validation and use *Is greater than* with a value of *10*. Use the message *Please enter 10 or less*.

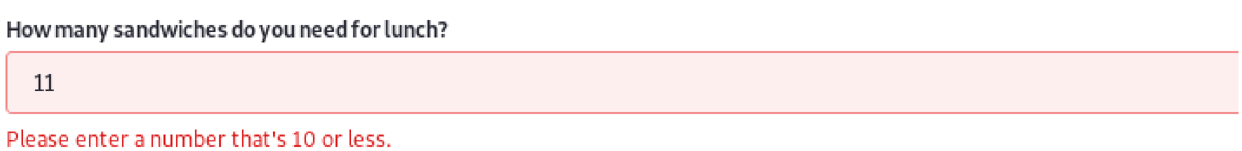

Figure 60.50: Make sure user-entered numeric data is within reasonable bounds. Nobody needs 11 sandwiches for lunch.

Note that numeric fields are text fields validated to allow only numeric data entry. That's why they're in the Customized Elements section of the form fields list. In addition, the property *My numeric type is* (can be Integer or Decimal) on the Basic tab of a numeric form is another form of validation.

60.21 Enabling CAPTCHA on Form Submissions

CAPTCHA prevents a bot from submitting forms. It's often used in login apps, but you can also use it in the Forms app.

To enable CAPTCHA, click the form's *Options* button () and select *Settings*. Enable the *Require CAPTCHA* setting, click *Done*, and save the form. That's all there is to it!

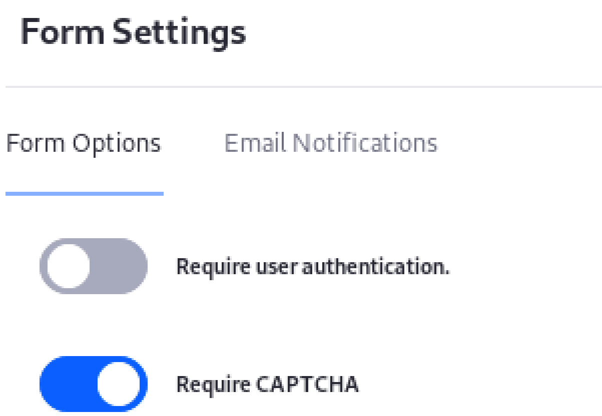

Figure 60.51: You can enable CAPTCHA for your form in the Form Settings window.

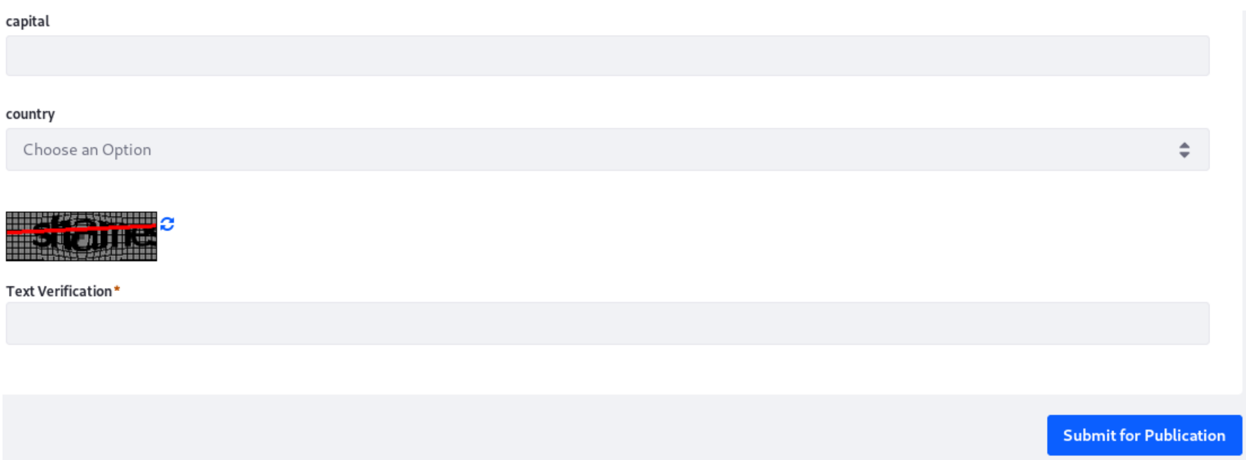

Figure 60.52: Once you enable CAPTCHA, your form has protection against bot submissions.

60.22 Form Notifications

You can configure the Forms app to send a notification email each time a form entry is submitted.

1. Open the form's *Form Settings* section by clicking the *Options* button () and selecting *Settings*.

2. Click the *Email Notifications* tab, enable the option *Send an email notification for each entry*, and fill out these fields:

   From Name: The sender's name. This could be the Site name, the form name, or anything else informative to the recipient.

   From Address: The sender's email address. You can use something like `noreply@example.com`, so that recipients don't try to reply.

   To Address: The recipient's email address (e.g., `test@example.com`).

   Subject: The email's subject. An informative subject line tells the recipient what happened. For example, An application for employment was submitted in The Lunar Resort*.

Note that if you enabled workflow for the form and it already sends a notification, you might not need to configure the Forms app to generate a notification.

60.23 Redirecting Users

When users submit a form, you can present them with another page indicating success or some other information related to their submission. Sometimes all you need is a success page, but other times you might want to send users to a specific URL.

Whatever your use case is, follow these steps to set up a redirect URL:

1. Open the form's *Form Settings* section by clicking the *Options* button (⋮) and select *Settings*.
2. Enter the redirect URL in the *Redirect URL* field.

That's it! Now when users submit the form, they're not left wondering what to do next.

60.24 Form Permissions

To access a form's permissions, first navigate to the Forms app in *Site Administration* (your site's menu) → *Content* → *Forms*. Then click the form's *Actions* button (⋮), and select *Permissions*.

By default, you can grant these permissions for a form:

Delete: Delete the form.

Permissions: Access and configure the form's permissions.

Add Form Instance Record: Submit form entries.

Update: Update form entries.

View: View the form.

Note that guest users can view and fill out forms by default. The *Guest* Role has *View* and *Add Form Instance Record* permissions.

Note: By default, all users inherit the Guest Role's permissions. The Guest Role represents unauthenticated visitors of your Site. If you want to let Guest users submit forms (the default setting), it makes sense that authenticated users can also. To disable automatic inheritance of the Guest Role's permissions, set this property in your `portal-ext.properties` file:

```
permissions.check.guest.enabled=false
```

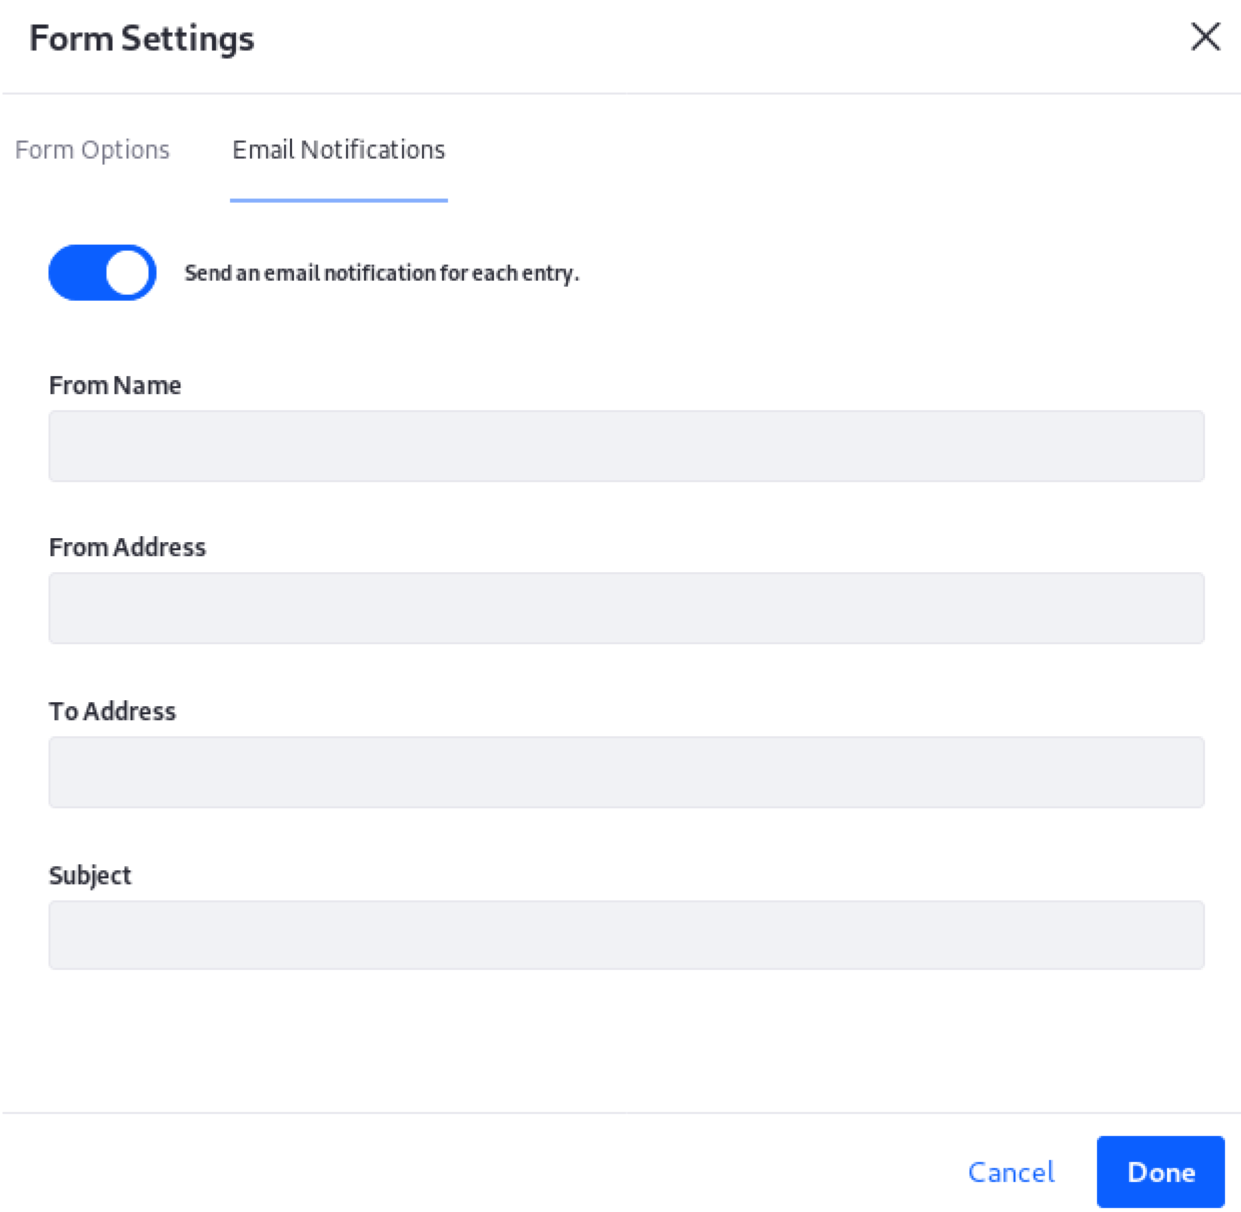

Figure 60.53: Configure email notifications each time a form entry is submitted.

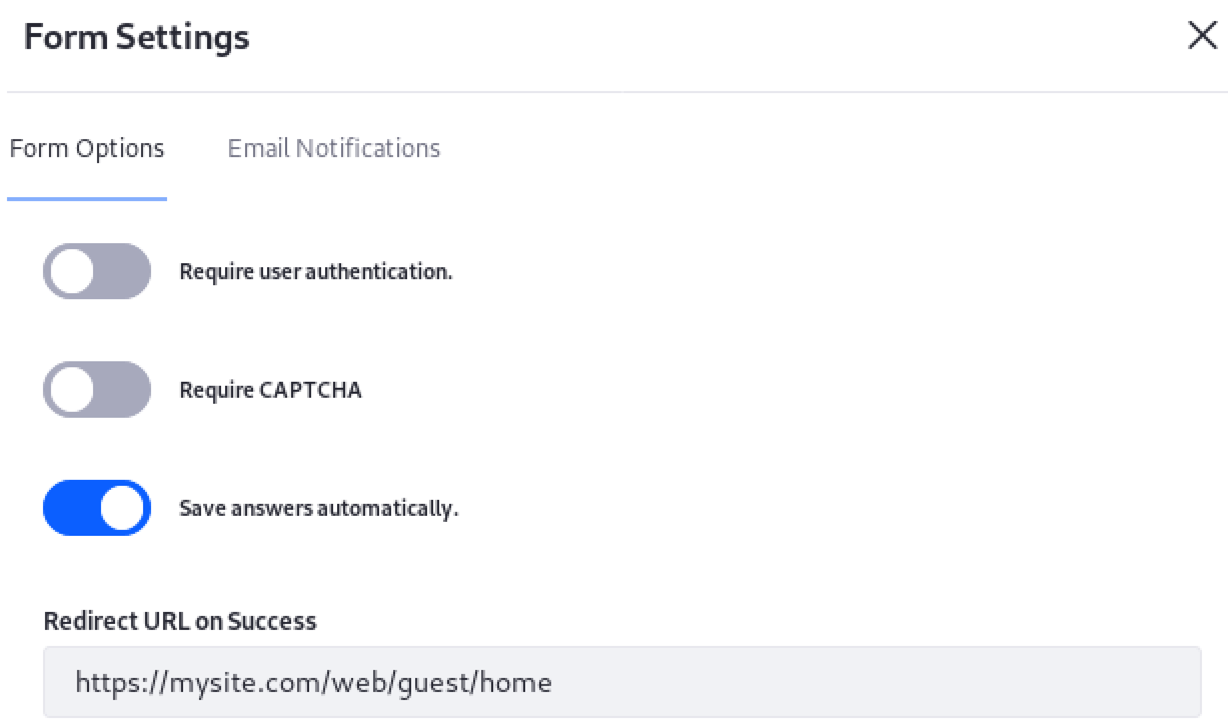

Figure 60.54: Redirect users after they submit a form.

60.25 Styling Form Pages

Let's face it: nobody likes an ugly, confusing form. Styling your form pages lets you make your forms user friendly. There are two features for styling your forms:

1. Create rows and columns for form fields.
2. Move fields from one location to another.

Sometimes it doesn't make sense to use the default single-column, vertically-oriented form. For example, a form with many fields can save space by putting them in different columns. You can also use a mixed approach, with each row broken into a different number of columns. The following screenshots show examples of these layouts.

By default, dragging a field onto a form page adds a field that occupies an entire row. Follow these steps to resize the field to make room for more fields in the row (columns):

1. Hover over the field to reveal its borders:
2. Drag the right or left edge of the field to resize it.
3. Add a new field to the remaining space in the row.

You can also move fields. To do so, follow these steps:

1. Hover over the field and the cursor becomes a hand.

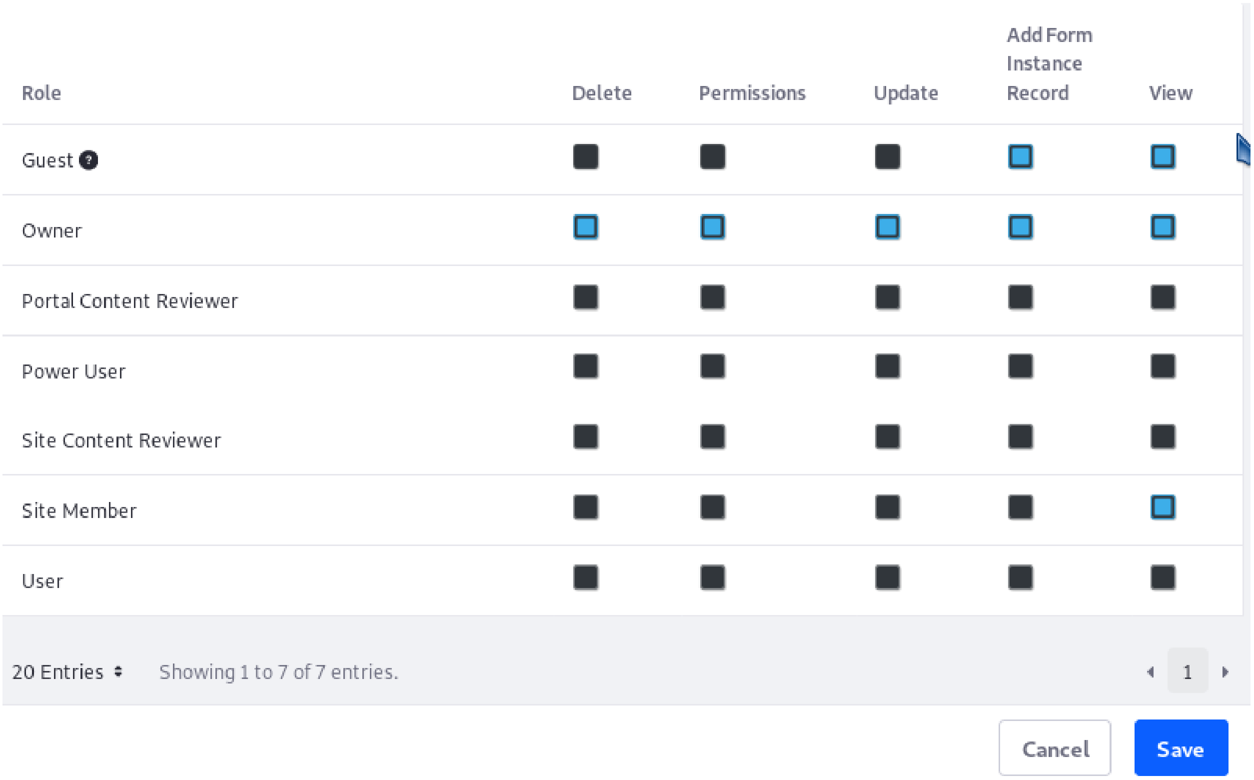

Figure 60.55: You can configure a form's permissions.

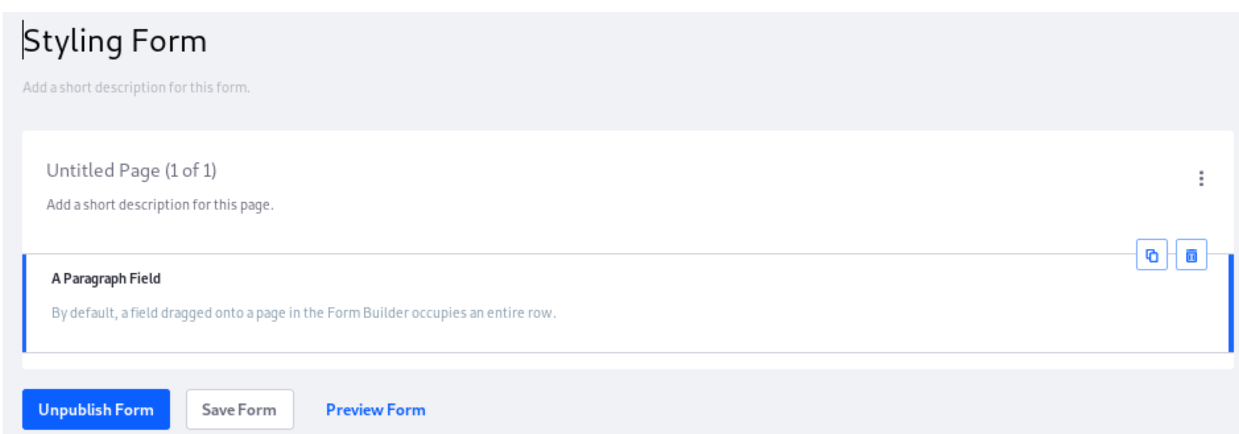

Figure 60.56: This is the default single-column, vertically-oriented form.

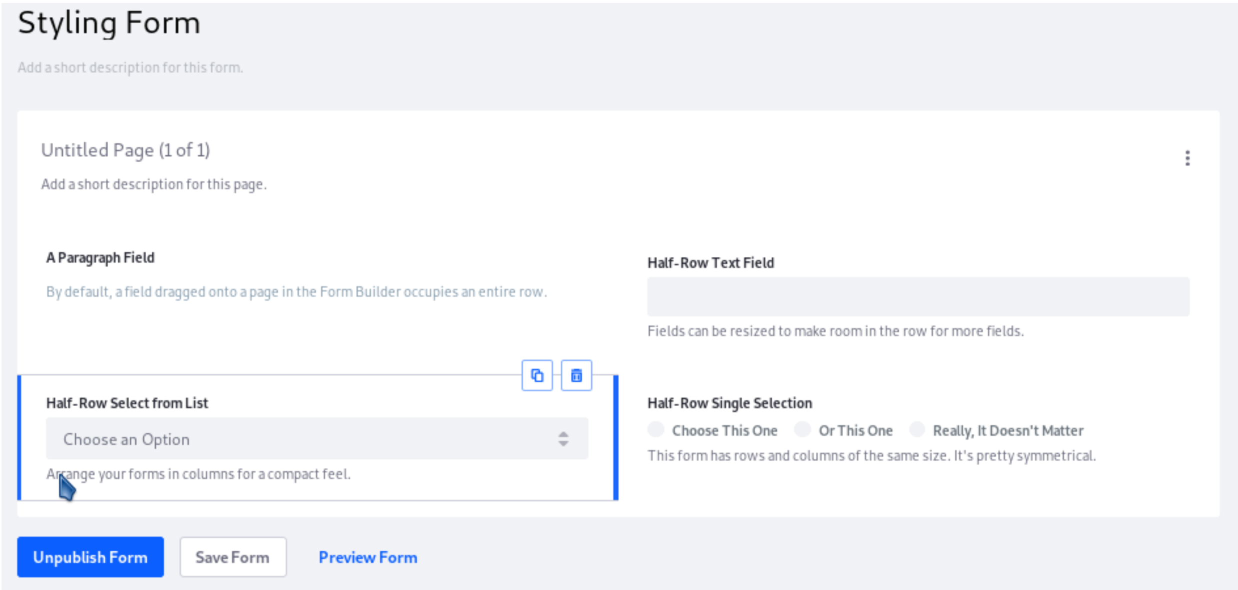

Figure 60.57: Putting form fields in multiple columns can give you more space.

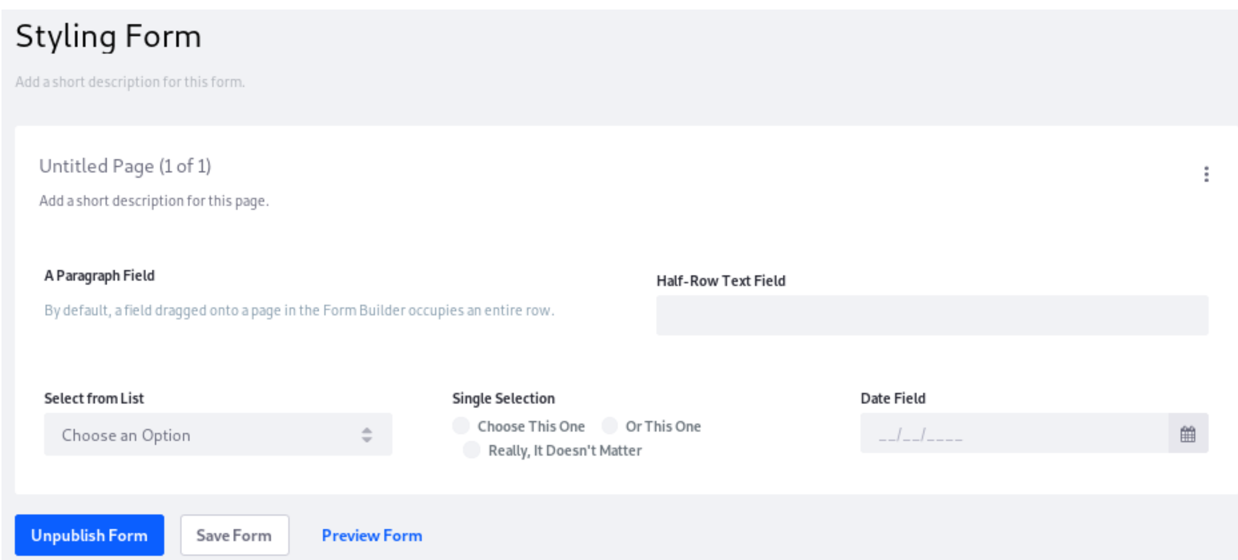

Figure 60.58: The first row is in two columns and the second row is in three columns.

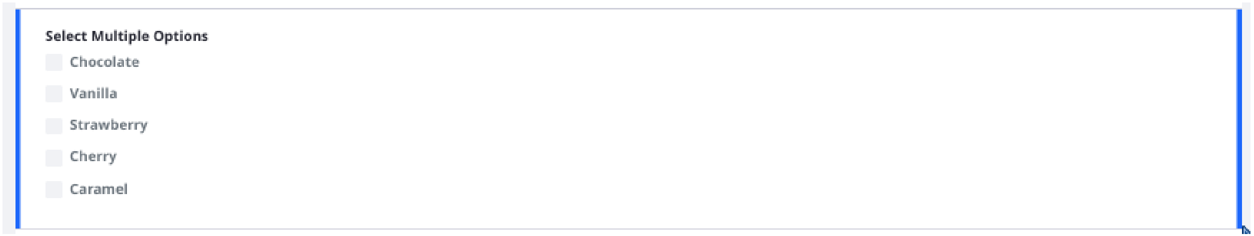

Figure 60.59: Form field borders.

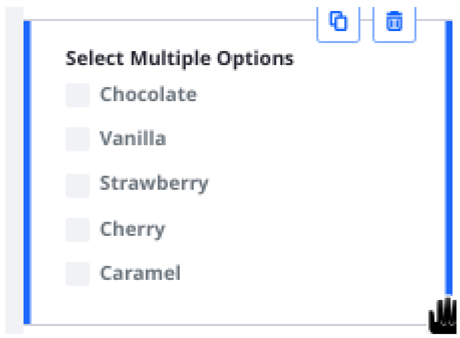

Figure 60.60: After resizing, the field is smaller.

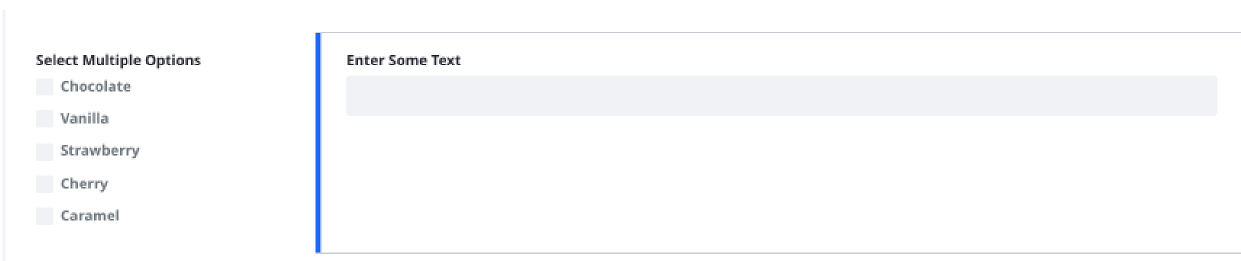

Figure 60.61: There are now two fields in the row.

2. Drag the field to an open location in the builder. Available locations are highlighted in blue and outlined with a dotted line. A field moved to a new row fills that entire row. A field moved to an existing row fills that row's remaining space.

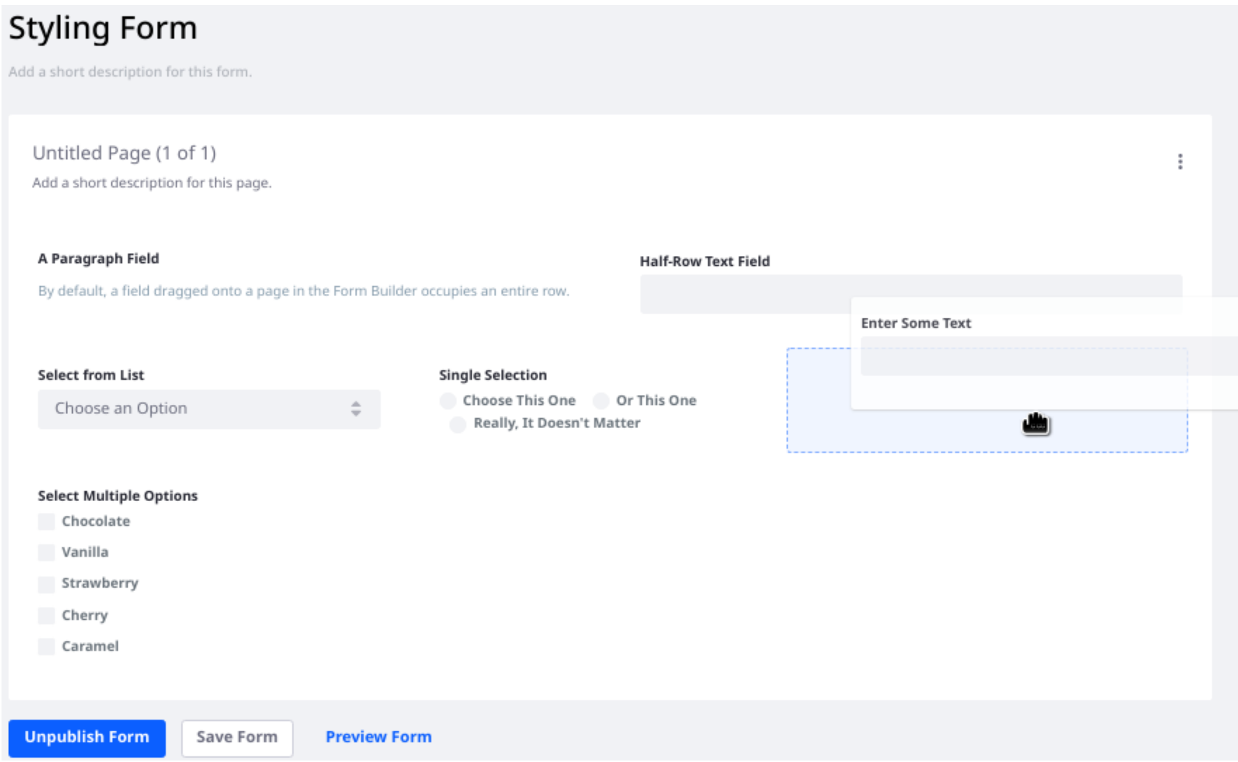

Figure 60.62: You can also move fields on form pages.

Chapter 61

Dynamic Data Lists

Dynamic data lists display forms created from field sets called *data definitions*. Data definitions consist of a form's field types (e.g., text, boolean, date, radio buttons, selector menus, etc.) and those fields' labels and settings. Data definitions effectively serve as data models for a dynamic data list. For example, you could create a data definition with two text fields: one for a user's name, and one for their comments. You could then display a form that gathers user feedback via a dynamic data list that uses that data definition.

To summarize:

- **Data Definitions:** Define a form's fields.
- **Dynamic Data Lists:** Display a form based on a data definition.

You can create one or multiple dynamic data lists from a single data definition. The user data entered for each dynamic data list is kept separate, even if the data definition is shared. For instance, you could use the example data definition above to create several dynamic data lists, and then place them anywhere you need to get feedback from users. Because each dynamic data list's form data is separate, you don't need to worry about trying to figure out which dynamic data list the user comment came from.

Dynamic data lists are flexible. You don't have to restrict dynamic data lists to simple input. You could create something as complex as an entire data entry system for real estate listings, or any other simple list-based application you can dream up.

You create data definitions and dynamic data lists in from the Site Menu's Content → Dynamic Data Lists application. Creating data definitions and lists doesn't require any coding. However, additional formatting can be added with FreeMarker templates.

These tutorials show you how to create and use data definitions and dynamic data lists:

- Creating data definitions.
- Creating dynamic data lists.
- Creating form and display templates.

61.1 System Configuration

There are two Dynamic Data Lists entries in System Settings. The Dynamic Data Lists Service entry contains one setting:

Add Default Structures This is enabled by default and pre-loads several embedded data definitions to base data lists on. Once loaded on portal startup, these definitions must be deleted manually from the Site Menu → Dynamic Data Lists application. This setting applies to the first start of a virtual instance.

The Dynamic Data Lists entry contains three settings:

Changeable Default Language If enabled, the default language of a data definition becomes changeable.

CSV Export Choose whether DDL records can be exported in CSV format with or without a warning, or disable this option. Here's what the warning says:

Warning: This CSV file contains user supplied inputs. Opening a CSV file in a spreadsheet program may be dangerous.

Default Display View Choose whether to use a table based default view or a list based default view.

61.2 Creating Data Definitions

Follow these steps to create a data definition:

1. Open the Menu () and expand your site's menu (the Site Administration menu). Then select *Content* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.

2. Click the *Options* button at the top-right () and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists any existing data definitions. Several definitions are embedded for common use cases like contacts, events, inventory, and more.

3. Click the *Add* button () to begin creating a new data definition. This opens the *New Data Definition* form.

4. Give your data definition a name. Note that the definition's name appears for any users filling out a dynamic data list that uses the definition. Then expand the *Details* section of the form and give it a description.

5. The Details section of the form also contains the field *Parent Data Definition*. This optional field lets you select an existing data definition (the parent) to form the basis of the new one (the child). The child definition inherits the parent's fields and settings, which you can then customize. When you create a dynamic data list from a child definition, it includes the fields of the parent and child definitions. This lets you use a common definition (the parent) as the basis of a specialized definition (the child). For example, if you were planning a rock climbing trip, you could use the default Events definition as the parent of a Rock Climbing Trip

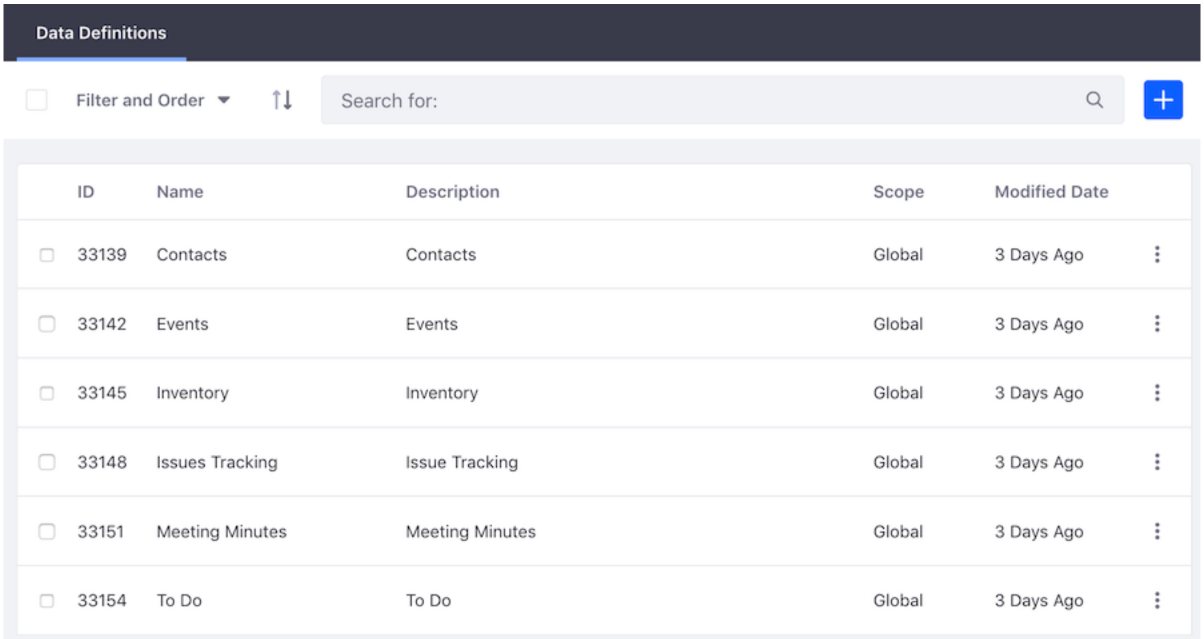

Figure 61.1: The Data Definitions screen.

definition that contains fields unique to rock climbing (e.g., climbing equipment availability, altitude, etc.).

To choose a parent definition, click the *Select* button below the *Parent Data Definition* field and then select an existing definition in the dialog that appears.

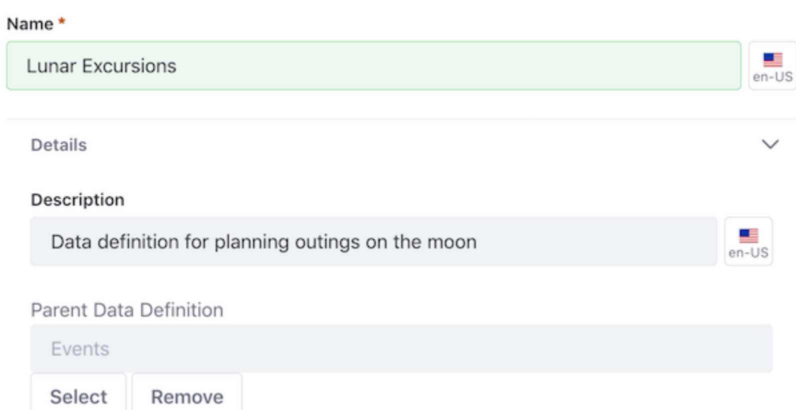

Figure 61.2: After naming your data definition, expand the Details section of the form and give your definition a description and parent definition, if desired.

6. Add the data definition's fields in the data definition designer, below the form's Details section. The designer's default *View* tab lets you create the definition in a WYSIWYG editor. You can click the *Source* tab to work with the definition's underlying JSON, but it's much easier to stick with the WYSIWYG editor.

   In the *View* tab select the *Fields* tab. Icons representing the field types are listed on one side and the data definition's canvas is on the other side. To add a field type to the definition, select its icon, drag, and drop it onto the canvas. By dragging a field onto a field that's already on the canvas, you can nest the new field in the existing field. When you mouse over a field on the canvas, the field action icons () appear. Clicking the + icon creates a duplicate of the current field and adds it below the current field. Clicking the trash can deletes the field.

The following fields are available:

- **Boolean:** A check box.
- **Color:** Specifies a color.
- **Date:** Enter a date. A valid date format is required for the date field, but you don't have to enter a date manually. When you select the date field a mini-calendar pops up which you can use to select a date.
- **Decimal:** Enter a decimal number. The value is persisted as a `double`.
- **Documents and Media:** Select a file from a Documents and Media library.
- **Geolocation:** Associate a location with the User's form entry.
- **HTML:** An area that uses a WYSIWYG editor to write and display HTML content.
- **Integer:** Enter an integer. The value is persisted as an `int`.
- **Link to Page:** Link to another page in the same site.
- **Number:** Enter a decimal number or an integer. The value is persisted either as a `double` or an `int`, depending on the input's type.
- **Radio:** Displays several clickable options. The default number of options is three but this is customizable. Only one option can be selected at a time.
- **Select:** This is just like the radio field except that the options are hidden and must be accessed from a drop-down menu.
- **Text:** Enter a single line of text.
- **Text Box:** This is just like the text field except you can enter multiple lines of text or separate paragraphs.
- **Web Content:** Select web content.

7. Edit field labels to reflect their intended data. A text field's default label is *Text*. To use the text field as a title, then you should change the field's label to *Title*. First select the field on the canvas. This automatically selects the *Settings* tab on the left. Alternatively, you can access the Settings tab by clicking the field's wrench icon. To edit a setting value, double-click it in the Settings table and enter the new value. The available settings are listed below.

   You can translate each of a data definition's field values to any supported locales. To specify a field value for a translation, select the flag that represents the locale and enter the field value for the locale.

   The following field settings are available. Note that some of these settings are only available for specific field types:

   - **Type:** The field's type (e.g., text, radio, etc.). This setting can't be edited, but a display template can reference it.
   - **Field Label:** The field's display name.
   - **Show Label:** Whether the field label is shown.
   - **Required:** Whether users must fill out the field (not available for Boolean fields).
   - **Name:** The field's internal identifier. You can use this value in a display template to read the field's data. This value is automatically generated, but you can change it if you wish.

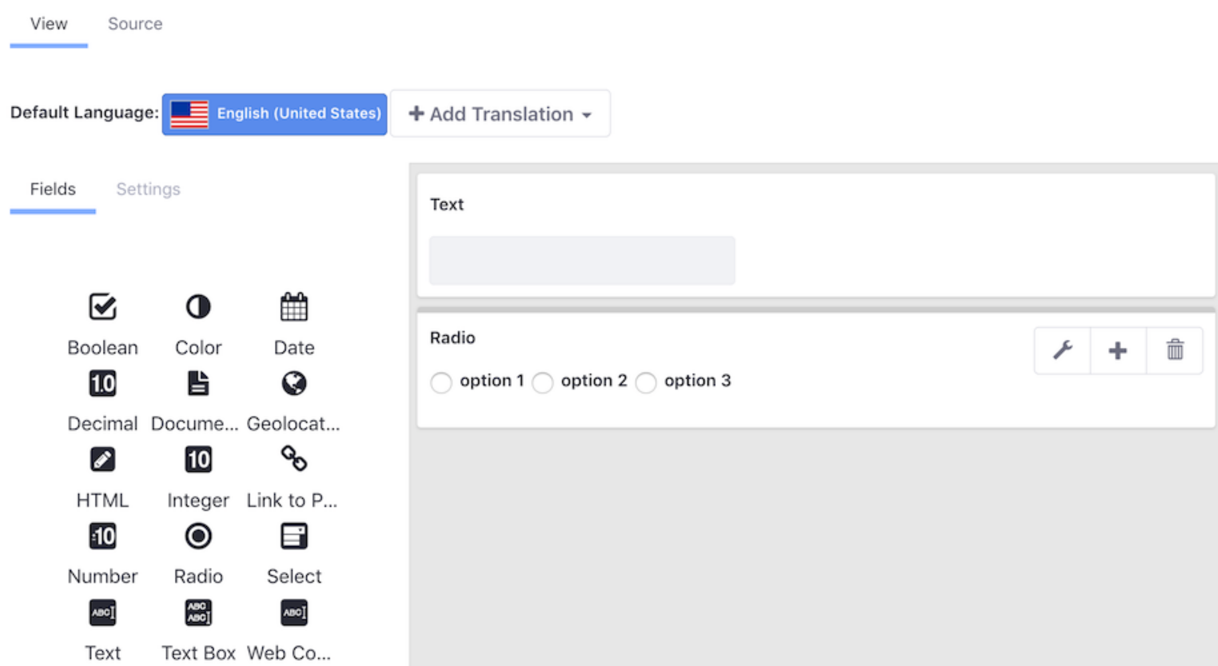

Figure 61.3: Use the data definition designer to add fields to the data definition.

- **Predefined Value:** The field's default value.
- **Tip:** Text to display in a tooltip.
- **Indexable:** Whether the field is indexed for search.
- **Localizable:** Whether the field can be translated.
- **Repeatable:** Whether users can make copies of the field.
- **Multiple:** Whether the user can select more than one option. This is only available for Select fields.
- **Options:** The options available for selection in Radio and Select fields. You can add and remove options, and edit each option's display name and value.

8. Click *Save* when you're done. Your new data definition then appears in the table with the pre-defined ones and any you've already added.

61.3 Managing Data Definitions

There are several ways to manage your data definitions. Of course, you can edit a data definition, but you can also configure its permissions, manage its templates, copy it, or delete it.

Follow these steps to access your data definitions:

1. Open the Menu ( ) and expand your site's menu (the Site Administration menu). Then select *Content* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.

2. Click the *Options* button at the top-right ( ) and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists the data definitions.

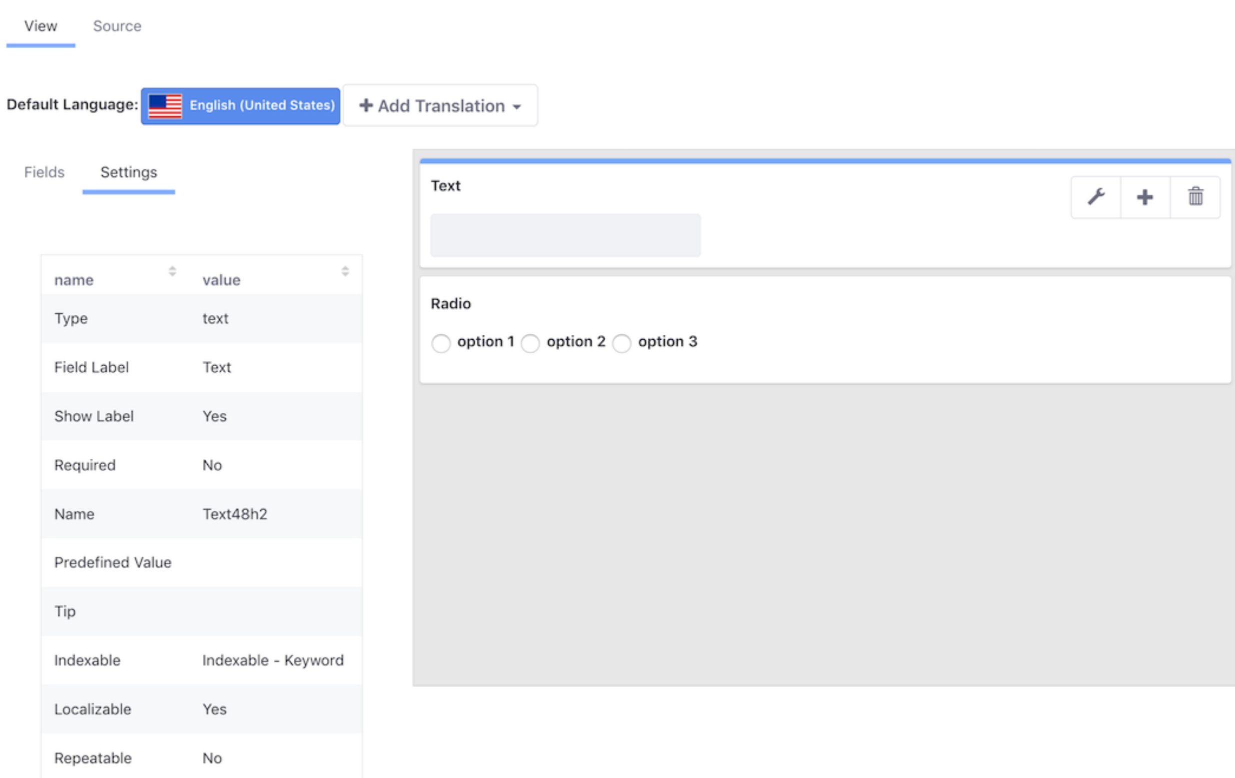

Figure 61.4: Configure the settings for each field in your data definition.

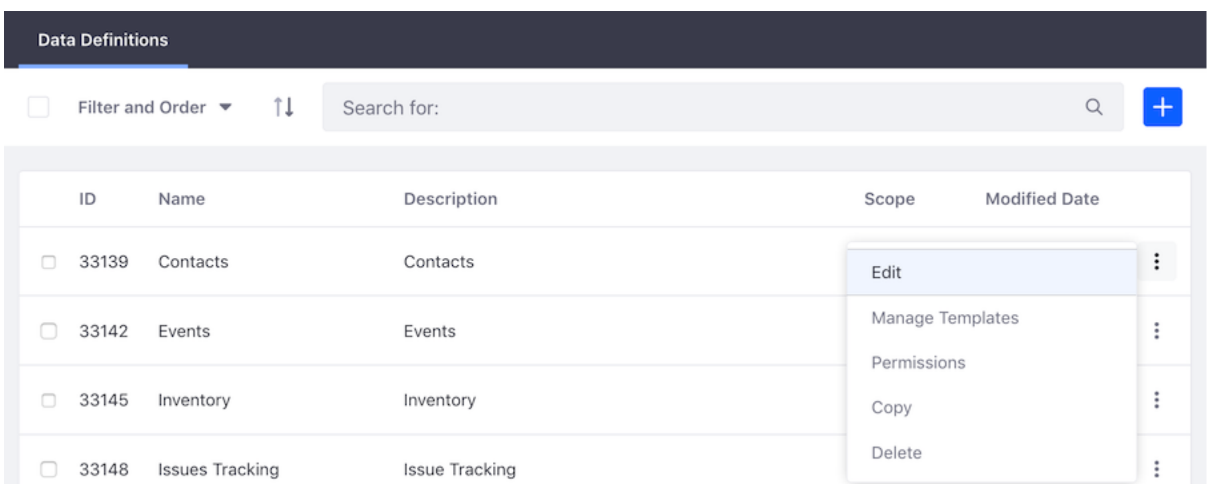

Figure 61.5: You can copy an existing data definition, manage its templates, and more.

You can manage your data definitions via the *Actions* menu (⋮) for each definition:

Edit: Edit the data definition. The edit screen uses the same form for creating data definitions. Note that if you edit a data definition referenced elsewhere (e.g., by a dynamic data list or display template), then you must update that reference.

Manage Templates: The *Manage Templates* screen creates and manages templates for the data definition. For details, see Using Templates to Display Forms and Lists.

Permissions: Configure the data definition's permissions. Note that these permissions are for an individual definition accessed through the Dynamic Data Lists application in *Site Administration → Content → Dynamic Data Lists*. For example, if Site members have View permission for a data definition, any Site member who also has a Role that can access the Dynamic Data Lists app and its data definitions can see this definition listed in the Manage Data Definitions screen. If you don't want this, remove the View permission for Site Member, and Site members won't see your data definition listed with the others.

Copy: The *Copy Data Definition* form copies the definition and its templates. In the form, give the copied definition a new name and description and select whether to also copy the original definition's templates. Click *Copy* when you're done. The copied definition then appears in the Data Definitions table with existing definitions. You can create new definitions based on existing ones, and then modify the copied one to suit your needs. You can, of course, edit any definition in the portal, but if you copy a definition instead, you can still access the original.

Delete: Delete the definition.

61.4 Creating Data Lists

There are two places to create dynamic data lists:

1. **Site Administration:** Open the Menu () and expand your Site's menu (the Site Administration menu). Then select *Content → Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table contains any existing lists. Click the *Add* button () to open *New List* form.

   To add Dynamic Data Lists, you must have permission to access the Dynamic Data Lists app in Site Administration.

2. **Dynamic Data Lists Display widget:** Navigate to the Site page where you want this widget and add it to the page from *Add* () → *Widgets* → *Collaboration* → *Dynamic Data Lists Display*. Then click the widget's *Add List* button. This opens the *New List* form.

   To do this, you must have permission to create a new list in the widget.

Either option leads to the New List form:

1. Give the list a name and a description.

2. Select the list's data definition: click *Select* under the *Data Definition* field, then click the definition you want to use.

3. To use a workflow with this list, select it from the *Workflow* field.

4. To change the list's default permissions, expand the form's *Permissions* section and make your selections.

5. Click *Save* to finish creating the list. Your new list appears in the table.

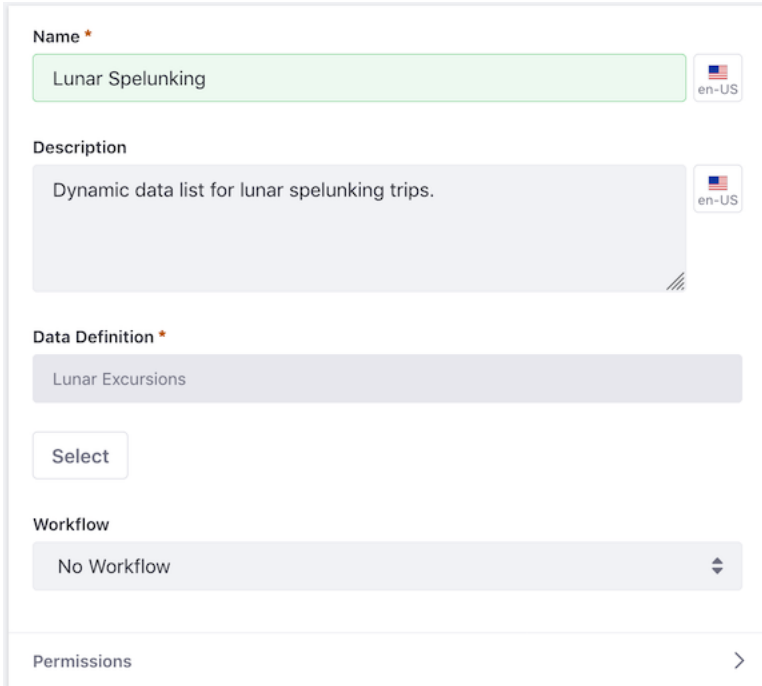

Figure 61.6: The New List form.

Creating List Records

By default, only administrators have permission to create list records. Follow these steps to give other users this permission:

1. Navigate to *Content* → *Dynamic Data Lists* in Site Administration.
2. Click *Actions* ( ) → *Permissions* for the list getting the new permissions.
3. Select *Add Record* for the Roles that should have that permission, then click *Save*. Allow unauthenticated Users to add records by giving Guest the Add Record permission.

Create new records in a list from the same places you can create the lists themselves:

1. **Site Administration:** In Site Administration, navigate to *Content* → *Dynamic Data Lists*. Click a list in the table to view any existing records, then click the *Add* button ( ). This opens a form based on the list's data definition, which you can then fill out and submit to create a new record. To do this, you must have permission to access the Dynamic Data Lists app in Site Administration.
2. **Dynamic Data Lists Display widget:** See above for instructions on adding this widget to a page. You must then configure the widget to display a list's records.

   To configure the widget to display a list's records:

   - Click the widget's *Select List* button.
   - In the dialog that appears, select a list, click *Save*, then close the dialog. The widget then displays the list's existing records.

To add a record:

- Click the widget's *Add* button ().
- Fill out the form that appears and click *Publish*.

See the section below for more information on configuring the widget.

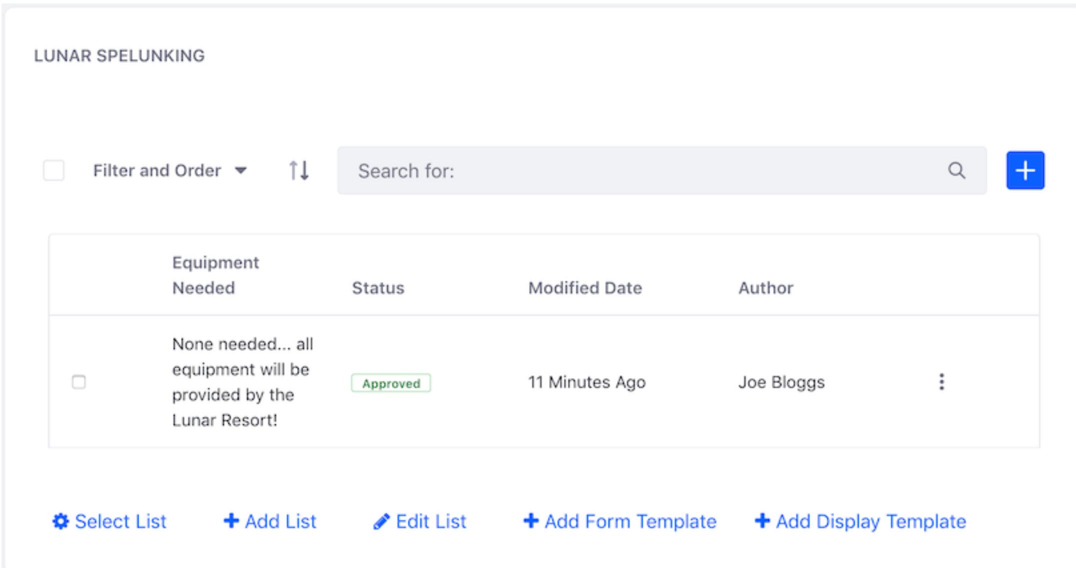

Figure 61.7: Dynamic Data Lists Display widget.

Configuring the Dynamic Data Lists Display Widget

The widget's default display template isn't exciting, but it shows the list's contents, and with permission, add and/or edit list items. To configure the widget, click its *Options* menu ( ) and select *Configuration*. This opens the Configuration dialog, with the *Setup* tab selected by default. The Setup tab contains two other tabs:

Lists: Select the list that the widget displays. The currently displayed list appears at the top of the tab, while the available lists appear in a table. To change the widget's list, select the list from the table and click *Save*.

Optional Configuration:

Display Template: Select the display template for the list.

Form Template: Select the form template for the list.

Editable: Whether users can add records to the widget's list.

Form View: Whether to display the Add Record form by default, instead of the List View. Note that even without this selected, users can still add records via the widget's *Add* button ().

Spreadsheet View: Whether the List View displays each record in a row, with columns for the record attributes.

When finished, click *Save* and close the Configuration dialog.

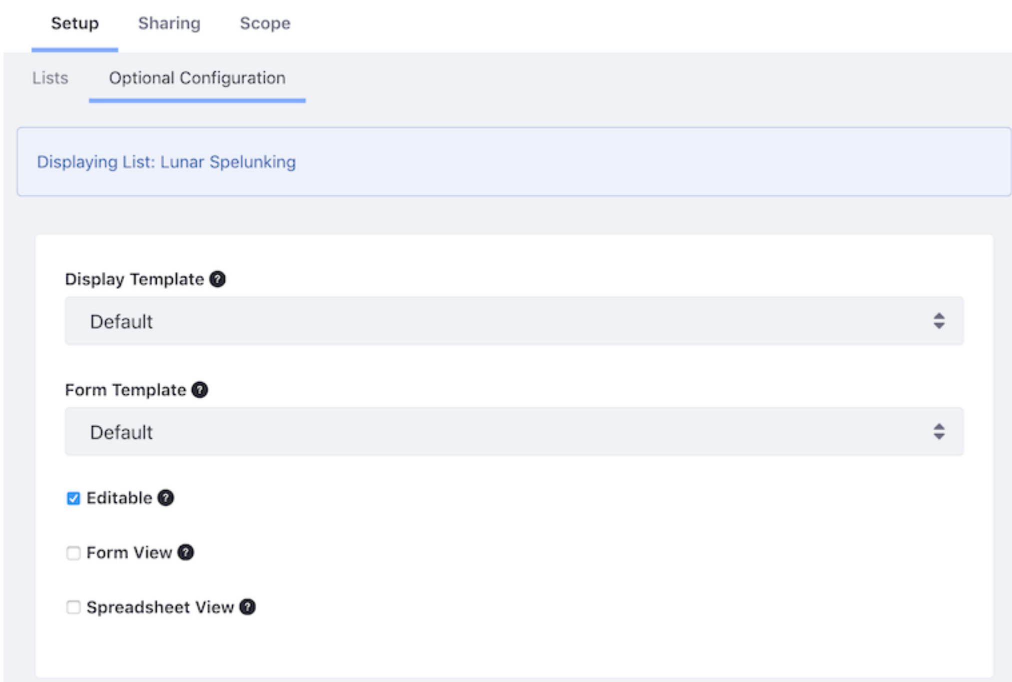

Figure 61.8: The Dynamic Data Lists Display widget's optional configuration.

61.5 Using Templates to Display Forms and Lists

After creating data definitions and lists, you can control how the form appears to users, and how the resulting list of records is displayed. You do this by creating templates for each view (form view for displaying the form and display view for the list of records) and selecting them in the DDL Display portlet. For example, you might need to create a form with a subset of a data definition's fields. Rather than creating a new definition, you can create a template that displays only the fields you want from the existing definition. You could also use a template to arrange fields differently, and/or with different labels and configuration options.

Data definitions can have as many form and display templates as you care to create (or none, if you're satisfied with the default templates). You then choose a list's template in the Dynamic Data Lists Display widget.

Managing Display and Form Templates

Since Display and Form Templates correspond to a particular data definition, they're accessed from the Data Definitions screen of the Dynamic Data Lists application in Site Administration. See the Creating Data Definitions article for instructions on accessing this screen.

The Data Definitions screen lists each definition in a table. To start working with a definition's templates, click the definition's Actions button (⋮) and select *Manage Templates*. This opens a screen that lists the definition's templates. You can edit, copy, delete, or configure permissions for a definition via its Actions button (⋮).

61.6 Creating Form Templates

Form templates control how the data entry form appears for a data definition. Follow these steps to create a form template for a definition:

1. Open the Menu () and expand your Site's menu (the Site Administration menu). Then select *Content* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.
2. Click the *Options* button at the top-right () and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists any existing data definitions.
3. Click the definition's *Actions* button ( ) and select *Manage Templates*. This lists the definition's templates.
4. Click the *Add* button () and select *Add Form Template*. This presents the same kind of graphical, drag-and-drop interface used to create definitions.
5. Give your template a name, then expand the *Details* section and give it a description. You can also use the *Mode* selector to select which mode the template applies to (*Create* or *Edit*).
6. Scroll down to the graphical designer in the *View* tab, and make your desired changes. For example, you can move or delete fields, change field labels, and more.
7. Click *Save* when you're finished.

Alternatively, you can create form templates from the Dynamic Data Lists Display widget:

1. Follow the instructions in the Creating Data Lists article for adding and configuring the widget in a Site page. Make sure to configure the widget to show the list you're creating a template for.
2. Click the widget's *Add Form Template* button. This opens the same form as above for creating a form template for the list's definition.

61.7 Creating Display Templates

For every data definition, you can create as many displays as you need. If you've created a form template that doesn't show all the fields of a particular data definition in the data list's form view, you probably don't want to display those fields in the list view, either. Modify the list view using Display Templates.

Note: If you're familiar with web content templates, display templates customize the display of a list in the same way. Display templates are written in FreeMarker or Velocity, pulling data from the data definition in the same way that web content templates pull data from their structures. Also similar to web content templates, display templates can be embedded in other display templates. This allows for reusable code, JavaScript library imports, or macros imported by Velocity or FreeMarker templates in the system. Embedding display templates provides a more efficient process when you have a multitude of similar data definitions. Just import an embedded display template and work off of it for your new display template.

As with form templates, you can create display templates from the Dynamic Data Lists app in Site Administration or the Dynamic Data Lists Display widget.

Follow these steps to create a display template from Site Administration:

1. Open the Menu (■) and expand your Site's menu (the Site Administration menu). Then select *Content → Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.

2. Click the *Options* button at the top-right (⋮) and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists existing data definitions.

3. Click the definition's *Actions* button (⋮) and select *Manage Templates*. This lists the definition's templates.

4. Click the *Add* button (+) and select *Add Display Template*.

5. Give the template a name, expand the *Details* section of the form, and give it a description. Also in the details section of the form, select the templating language to use from the *Language* selector. You can choose FreeMarker, or Velocity.

6. In the *Script* section of the form, create your template in the editor using the templating language you chose in the previous step. The palette to the left of the editor contains common variables. Click a variable to insert it in the editor. The editor also autocompletes. In a FreeMarker template, type `${`, which opens an autocomplete list of common variables. Select a variable to insert it in the editor. Alternatively, you can upload a complete script file via the *Browse* button below the editor.

7. Click *Save* when you're done creating the template.

Alternatively, you can use the Dynamic Data Lists Display widget to create display templates:

1. Follow the instructions in the Creating Data Lists article for adding and configuring the widget in a site page. Make sure to configure the widget to show the list's definition you're making into a template.

2. Click the widget's *Add Display Template* button. This opens the same form as above for creating a display template for the list's definition.

Display Template Editor

Helper variables are available in the template editor. These provide access to most of the data that you'll use in creating Display Templates. The variables under the heading Data List Variables let you inject specific information about the data definition the template is being created for:

Data Definition ID: `reserved_ddm_structure_id`
Data List Description: `reserved_record_set_description`
Data List ID: `reserved_record_set_id`
Data List Name: `reserved_record_set_name`

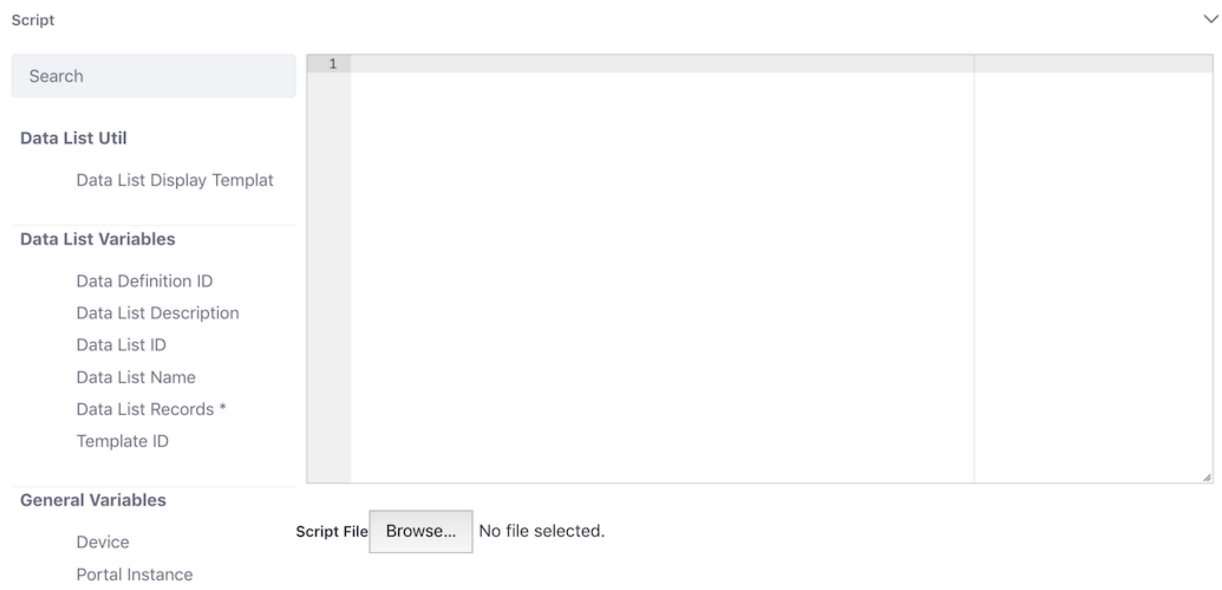

Figure 61.9: Create your display template in the editor.

Template ID: `reserved_ddm_template_id`

Inside a template, these variables give the ID for the record set as well as the name, description, and data definition.

Display the list of records by retrieving them and assigning them to the handy `records` variable. Retrieve the list's records from `DDLDisplayTemplateHelper`, which contains these functions:

`getDocumentLibraryPreviewURL`

`getHTMLContent`

`getLayoutFriendlyURL`

`getRecords`

`renderRecordFieldValue`

`DDLDisplayTemplateHelper` performs common tasks. Use the `getRecords` method to access a data definition's entries and assign them to a `records` variable:

```
<#assign records = ddlDisplayTemplateHelper.getRecords(reserved_record_set_id)>
```

This *fetches* the records of the associated data list. You haven't done anything with them yet, so your display is still empty. To list all the records, use the *Data List Records* helper in the sidebar of the template editor. Remember to place your cursor in the proper place in the template editor window, then click *Data List Records*. This code appears at the cursor:

```
<#if records?has_content>
    <#list records as cur_record>
        ${cur_record}
    </#list>
</#if>
```

This default code snippet spits out everything in the database for the given data definition, which is ugly and practically useless:


```
{uuid=52c4ac1c-afe7-963c-49c6-5279b7030a99, recordId=35926, groupId=20126,
companyId=20099, userId=20139, userName=Test Test, versionUserId=20139,
versionUserName=Test Test, createDate=2018-07-16 14:31:51.056,
modifiedDate=2018-07-16 14:31:51.058, DDMStorageId=35927, recordSetId=35922,
recordSetVersion=1.0, version=1.0, displayIndex=0, lastPublishDate=null}
```


Here's a simple example template that uses a list based on the embedded Contacts data definition, and only displays the Company Name and Email fields in a bulleted list:

```
<#assign records = ddlDisplayTemplateHelper.getRecords(reserved_record_set_id)>

    <h1>Here are contacts by company name and email address.</h1>

    <#if records?has_content>
        <#list records as cur_record>
            <ul>
                <li>
                    <#-- The below gets the Company field and wraps it in an <em> tag -->
                    Company Name: <em>${ddlDisplayTemplateHelper.renderRecordFieldValue(cur_record.getDDMFormFieldValues("company")?first, locale)}<
                    <#-- The below gets the Email field  and wraps it in an <em> tag -->
                    Email: ${ddlDisplayTemplateHelper.renderRecordFieldValue(cur_record.getDDMFormFieldValues("email")?first, locale)}
                </li>
            </ul>
        </#list>
    </#if>
```

Here's what it looks like:

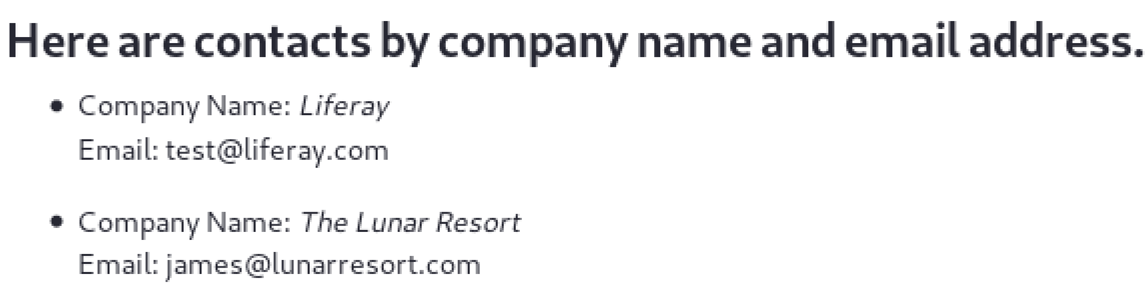

Figure 61.10: Extract appropriate display information, rather than spitting out the whole object.

Now you're prepared to make data lists beautiful using Display Templates.

Chapter 62

Workflow

Liferay's workflow engine is named *Kaleo*. In Greek, Kaleo means "called ones," which is appropriate for a workflow engine that calls users to participate in a review process designed for them. Kaleo workflow makes it possible to define any number of simple to complex business processes/workflows, deploy them, and manage them through a portal interface. The processes have knowledge of Users, Groups and Roles. You don't have to write a single line of code to accomplish this: it can be defined in an XML document, or if you're a Liferay Digital Experience Platform (DXP) customer, you have access to a nifty visual designer for workflows, called Kaleo Designer.

There are several steps to effective workflowing:

- Designing review processes in XML
- (DXP only) Designing review processes in Kaleo Designer
- Uploading workflow definitions
- Activating workflow for enabled assets
- Managing Workflow definitions
- Sending assets through review

After all that, you'll be familiar with using Liferay's Kaleo workflow to set up approval processes for any workflow-enabled content.

62.1 What's New with Workflow

There are some noteworthy enhancements to the workflow functionality:

- Duplicate a workflow.
- Preview and restore a prior version.
- Save draft workflows and retain them in an unpublished state (this is a new feature for Liferay Portal CE workflow users).

- All Control Panel workflow functionality is unified under one entry, found at Control Panel → Configuration → Workflow.
- Create workflows (using a text editor for Liferay Portal CE users, and the Workflow Designer for Liferay DXP customers) at Control Panel → Configuration → Workflow. Just click the Add button.

Create your own workflows from scratch or leverage existing workflows.

62.2 Embedded Workflows

In addition to the Single Approver definition, there are some workflow definitions that ship with Liferay DXP but are not pre-installed, since they're primarily included for test cases. They can be found in the Liferay source code in `/modules/apps/portal-workflow/portal-workflow-kaleo-runtime-impl/src/main/resources/META-INF/definitions`, or in your Liferay installation itself. Open your Liferay installation's `osgi/portal/com.liferay.portal.workflow.kaleo.runtime.impl.jar`, and then find and open the `com.liferay.workflow.kaleo.runtime.impl-[version].jar`. The definitions are in the `META-INF/definitions` folder (see the diagrams here).

Chapter 63

Kaleo Designer

With the proper permissions, users can publish assets. Even if your enterprise has the greatest employees in the world, many of the items they want to publish must still be reviewed, for a variety of reasons. The Kaleo Designer lets you design workflow definitions so your assets go through a review process before publication.

With Kaleo Designer, you develop workflow definitions using a convenient drag and drop user interface. You don't need to be familiar with writing XML definitions by hand. Some of the features can be enhanced, however, if you're familiar with Groovy, a supported Java-based scripting language. All that is to say, don't be scared off when you come to a block of code in these articles. Just decide if you need the feature and find someone familiar with Java or Groovy to help you out.

Note: By default, only one workflow definition is installed: the Single Approver Workflow definition. What you might not know is that you have access to several others too. Look in `[Liferay_Home]/osgi/portal/` and find the `com.liferay.portal.workflow.kaleo.runtime.impl-[version].jar`. Open it up and then look in `META-INF/definitions`. You'll see the following workflow definitions:

```
category-specific-definition.xml
legal-marketing-definition.xml
single-approver-definition.xml
single-approver-definition-scripted-assignment.xml
```

To work with any of these definitions in Kaleo Designer, extract them from the JAR file first. Once you have the XML files locally,

1. Add a new workflow. Go to Control Panel → Configuration → Workflow, and click the Add button (+).

2. Go to the Source tab.

3. Click *import a file* and upload the XML file.

4. Name the definition appropriately, and click either *Save* (to save it as a draft) or *Publish* (see below for more information on saving and publishing).

Now you can begin exploring or modifying the definition.

It's time to start exploring the Kaleo Designer and its features.

63.1 Managing Workflows with Kaleo Designer

The Kaleo Designer gives you an intuitive interface for creating workflow definitions, from the simplest approval processes to the most complex business processes you can imagine. It features a drag and drop interface, workflow definition versioning, and a graphical representation of definitions and their nodes. Without Kaleo Designer, you'd have to hand-craft your entire workflow definition in XML. With Kaleo Designer, you might never need to look at a single line of XML. Of course, the Kaleo Designer can also directly manipulate the XML (using the *Source* tab) if you find it convenient.

Adding New Workflow Definitions with Kaleo Designer

Access Kaleo Designer by going to the Control Panel, then clicking Configuration → Workflow. Click the Add icon ().

Give the workflow definition a title and you're ready to start designing your workflow.

Saving and Publishing Workflow Definitions

First, look below the canvas to see the buttons that let you *Save* or *Publish*. Saving the definition as a draft lets you save your work so it's not lost (due to a timeout, for example). It won't be published (and assignable to assets), and it won't be considered a version until the Publish button is clicked. Each time you save the workflow as a draft, a new revision is added to the Revision history. To see the Revision history and manage workflow versions, open the Info sidebar () and click *Revision History*.

Adding Nodes

A new workflow is already populated with a start node, an end node, and a transition between them. To make the workflow the way you want it, add nodes to the workflow.

1. *Drag* a node from the *Nodes* palette on the right of the designer and *drop* it on the canvas.

2. You'll see it's not connected to other nodes by a transition, so right now it can't be used in the workflow. Delete the existing transition and then you can make new transitions to direct the *flow* of your workflow (see more about transitions below if you're not sure what they're for or how to use them in Kaleo Designer).

Alternatively, start by deleting the default transition, then click the edge of the start node, drag a new transition from the start node to a blank spot on the canvas, and release it. You're prompted to create a node at that spot, because you can't have a transition without a starting point and an ending point on a node.

That's it. Of course, if you drag, say, a *Task* node onto the canvas, it must be configured.

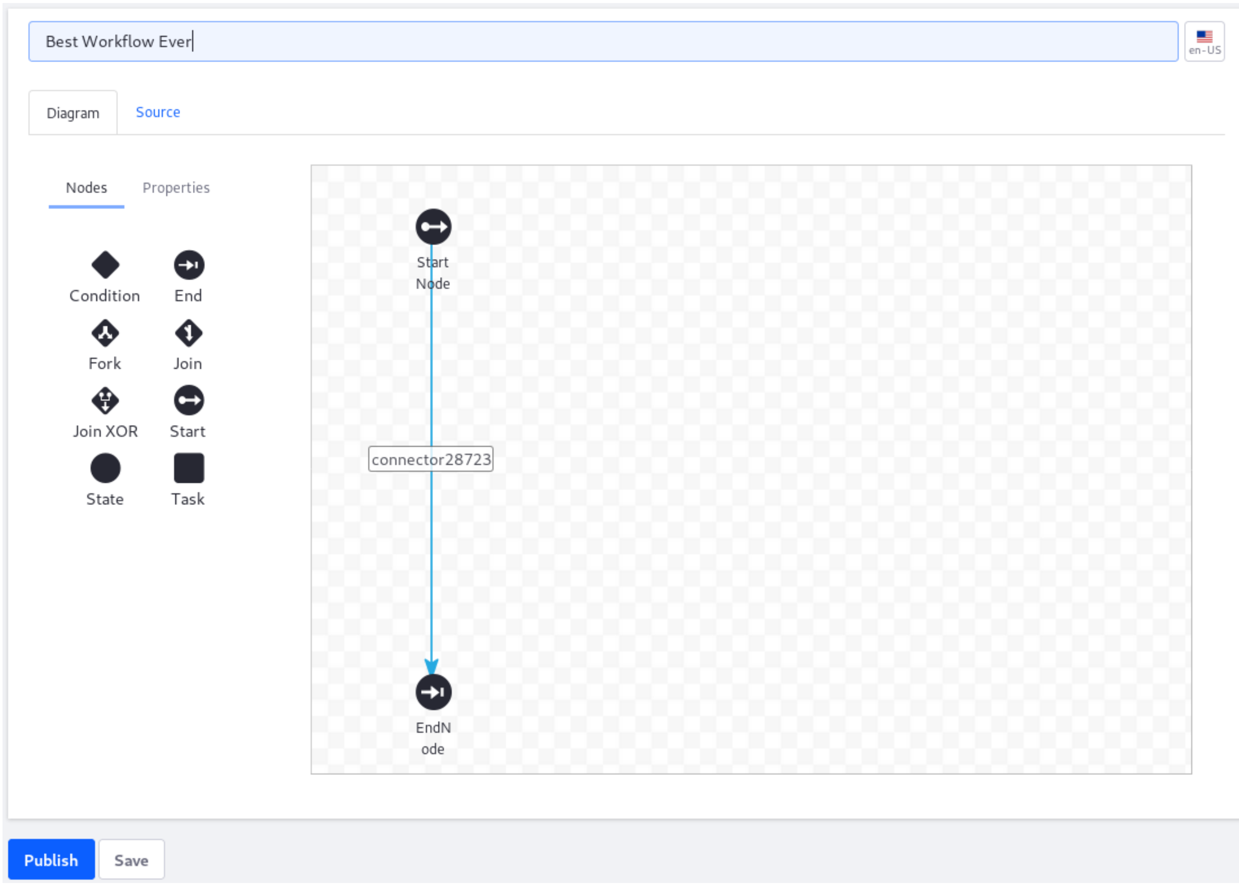

Figure 63.1: The Kaleo Designer's graphical interface makes designing workflows intuitive.

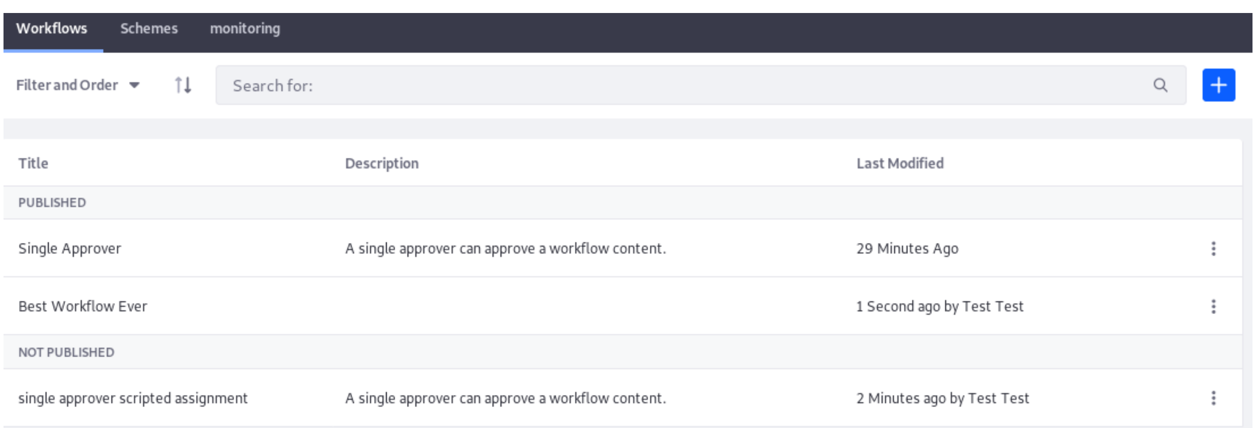

Figure 63.2: View a list of the current workflows that can be edited in the Kaleo Designer.

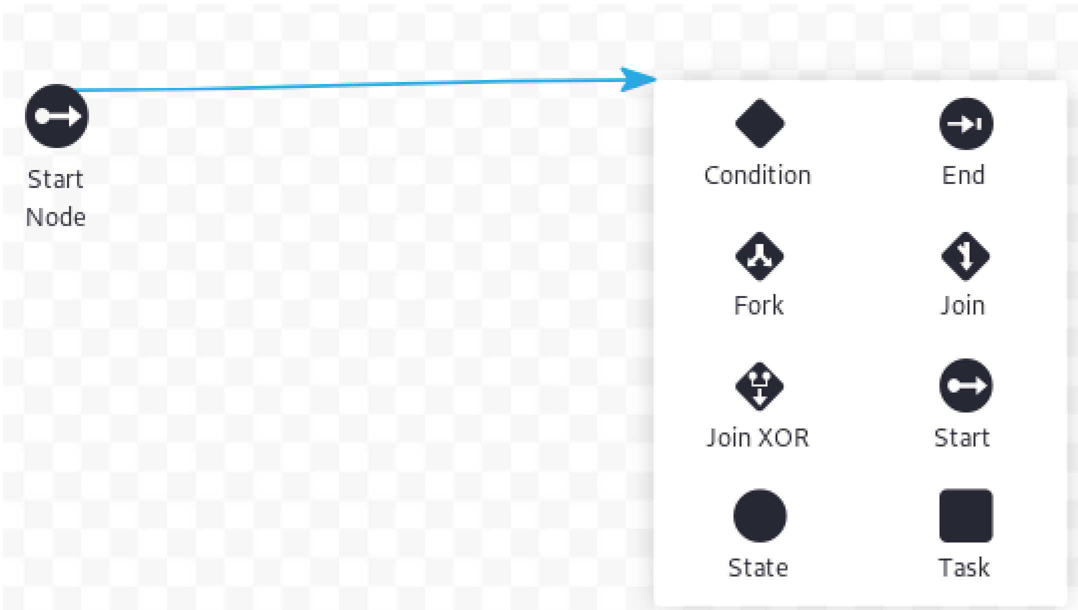

Figure 63.3: You can add a node by creating a transition that ends at a blank spot on your Designer canvas.

Node Settings

Now you know how to add nodes to the workflow definition. By default you have three things added to your canvas: a start node, a transition, and an end node. Think of the *EndNode* as the point in the workflow where an asset reaches the *Approved* status. The *StartNode* is where the asset goes from the *Draft* status to *Pending*. You might decide to name your nodes to reflect what's happening in each one. To name a node, double click it, and its *Settings* appear. Then double click the value of the *Name* property and you can edit the name. Click *Save* when you're done.

Of course, there's more you can do besides changing node names. Actions, Notifications, and Assignments can be used to make your workflow definition useful and interactive. Keep reading to learn about these features.

63.2 Workflow Definition Nodes

Once you know the basics of creating workflow definitions with the workflow Designer, you can get into the details. In this tutorial you'll learn about Actions and Notifications, two important features your workflow nodes can use. You'll also learn how to affect the processing of the workflow using Transitions, Forks, Joins, and Conditions.

There are several node types you can use in workflow definitions:

- Task nodes
- Fork and Join nodes
- Condition nodes
- Start nodes
- End nodes
- State nodes

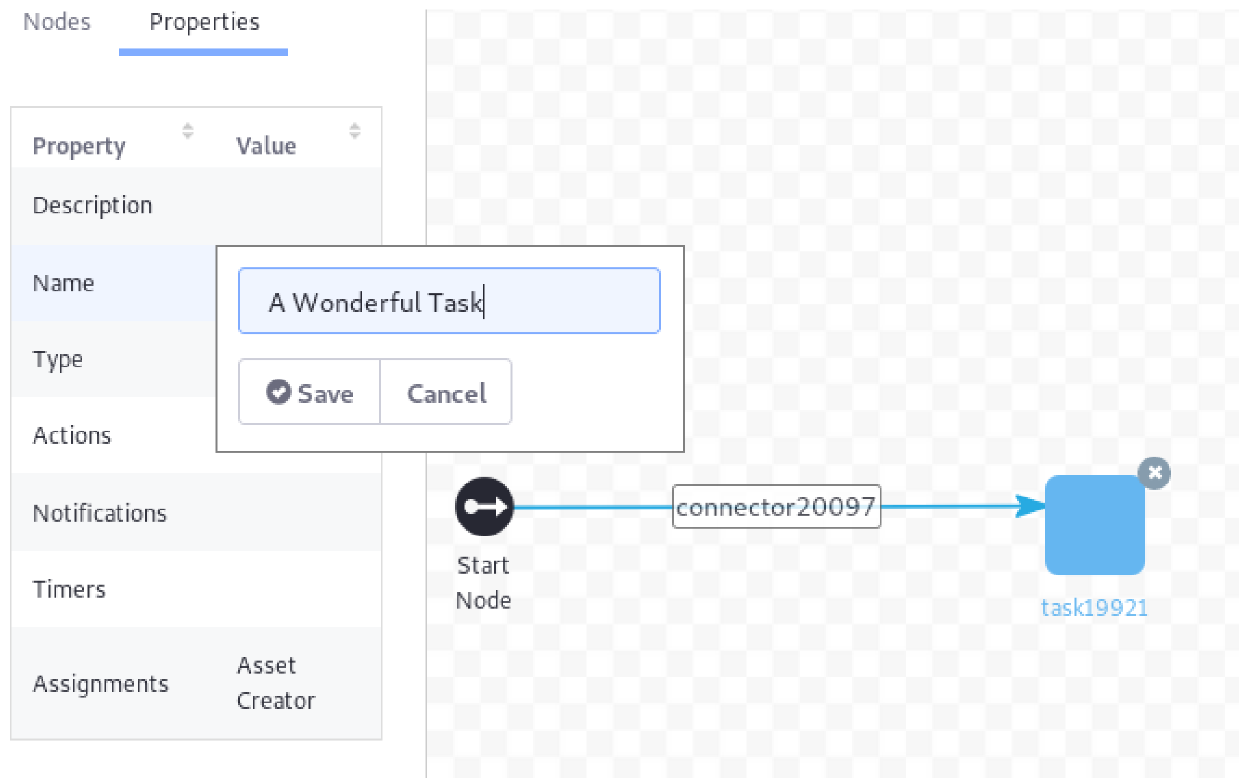

Figure 63.4: You can edit a node's settings.

Because they're the most complex node, and often the meat of your workflow definitions, Task nodes are covered separately.

Fork, Join, and Condition nodes are discussed, along with Transitions, in a tutorial on workflow processing, since they're used for affecting the processing of the workflow.

This tutorial discusses State nodes, Start nodes, and End nodes, along with Actions and Notifications.

Node Actions and Notifications

Any node can have Actions and Notifications.

Actions

Actions do additional processing before entering the node, after exiting a node, or once a task node is assigned. They're configured by accessing a node's Properties tab, then double clicking *Actions*.

The Single Approver workflow contains an Update task with an action written in Groovy that sets the status of the asset as *denied*, then sets it to *pending*.

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("denied"), workflowContext);
WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("pending"), workflowContext);
```

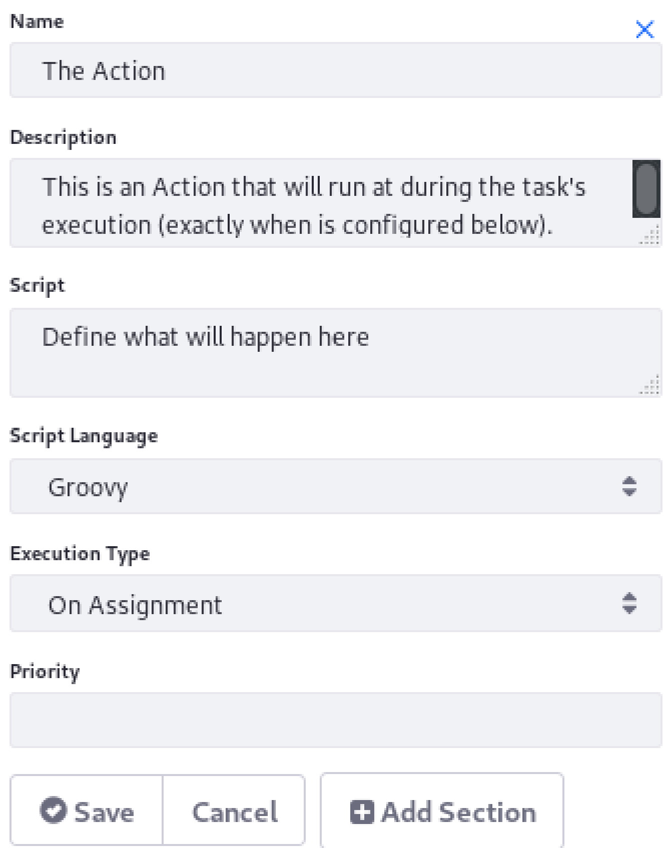

Figure 63.5: You can add an Action to a Task node.

Why would the action script first set the status to one thing and then to another like that? Because for some assets, the *denied* status sends the asset creator an email notification that the item has been denied.

The end node in your workflow definition has an action configured on it by default, on entry to the end node:

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("approved"), workflowContext);
```

This is a Groovy script that updates the status to *approved*, since that's usually the goal of a workflow process.

You can do something simple like the actions above, or you can be as creative as you'd like.

It's good to assign a task to a user, and it's even more useful if the user can get notified of a workflow task.

Notifications

Notifications are sent to tell task assignees that the workflow needs attention or to update asset creators on the status of the process. They can be sent for tasks or any other type of node in the workflow. To set up notifications, double click on *Notifications* in a node's Properties tab and create a notification.

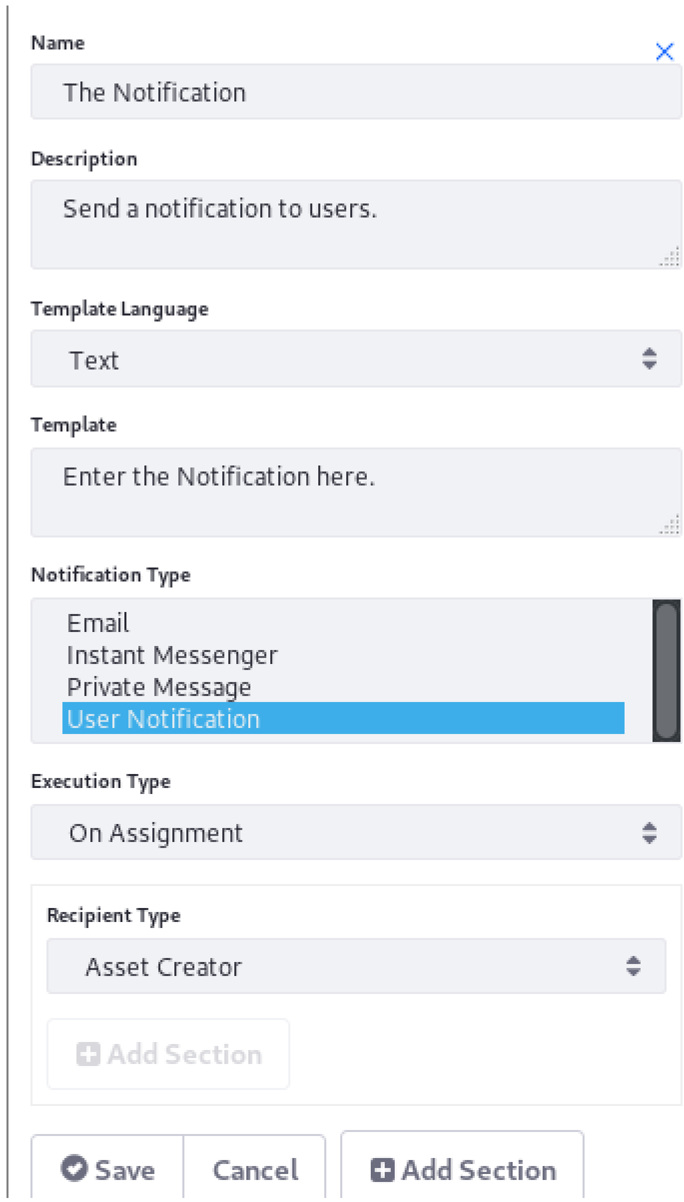

Figure 63.6: You can send a Notification from a Task node.

You must specify the Notification Type, and you can choose User Notification or Email. You can use Freemarker or Velocity if you need a template, or you can choose to write a plain text message.

Note: Instant Messenger and Private Message also appear as Notification Type options, but these are non-functional and will be removed in a future version.

Here's a basic Freemarker template that reports the name of the asset creator and the type of asset in the notification:

```
${userName} sent you a ${entryType} for review in the workflow.
```

You can also choose to link the sending of the notification to entry into the node (On Entry), when a task is assigned (On Assignment), or when the workflow processing is leaving a node (On Exit). You can configure multiple notifications on a node.

Commonly, the assignment and notification settings are teamed up so a user receives a notification when assigned a task in the workflow. To do this, choose *Task Assignees* under Recipient Type when configuring the notification.

Note: The *from name* and *from address* of an email notification are configurable via portal properties. Place these settings into a `portal-ext.properties` file, in your Liferay Home folder. Then restart the server:

```
workflow.email.from.name=
workflow.email.from.address=
```

These can also be set programmatically into the `WorkflowContext`, and the programmatic setting always takes precedence over the system scoped portal property.

Start and End Nodes

Start and end nodes kick off the workflow processing and bring the asset to its final, approved state. Often you can use the default start and end nodes without modification. If you want to do some more processing (in the case of a start node), add an action to the node using the Properties tab as described in the section on Actions above.

End nodes have a default action that sets the workflow status to Approved using the Groovy scripting language:

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("approved"), workflowContext);
```

Feel free to add more to the action script if you need to do additional processing.

By default, there's a transition connecting the start node and end node, but you'll probably want to delete it, since most workflows don't proceed straight from the initial state to approved.

State Nodes

State nodes can have Notifications and Actions. The default end node added by workflow Designer is a pre-configured state node that sets the workflow status to Approved. Perhaps you want to create a node that sets the status to *Expired*. You could create a state node for it by dragging one onto your workflow Designer canvas, then configuring an action in it that sets the status to Expired. Here's what it would look like in Groovy:

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("expired"), workflowContext);
```

Next, you'll learn to do parallel processing using fork and join nodes.

63.3 Affecting the Processing of Workflow Definitions

Workflow definitions all contain nodes: a Start Node, and End node, and at least one Task node. You might know that for the workflow to progress from one node to the other, you need Transitions. In this tutorial you'll learn about using transitions to move the asset through the workflow from node to node, but you'll also learn about some other features that affect the processing of the workflow.

- Transitions
- Forks
- Joins
- Conditions

Start by learning about the ever important Transition.

Transitions

What are transitions? Workflow transitions connect one node to another. On exiting the first node, processing continues to the node pointed to by the transition. Every time you create an arrow from one node to another, Workflow Designer creates a transition.

Each node you add has a pop-up menu letting you delete the node. As you hover your mouse over the edges of a node, notice your mouse pointer changes to a cross. The cross indicates you can connect the current node to another node. Hold down your mouse button and drag the mouse to start drawing your transition to another node. If you stop before reaching the edge of the next node, a pop-up displays node types you can create and connect to on-the-fly. To connect with an existing node, continue dragging the connector to that node.

When developing workflows in Kaleo Designer, make sure you go through all the transitions and name them appropriately. By default, these transitions get system generated names, so rename them to something more human-readable, as they're displayed to workflow users as links that send the item to the next step in the workflow.

To rename transitions, click on the arrow representing the transition and use the Properties tab to set the name just like you do for a node.

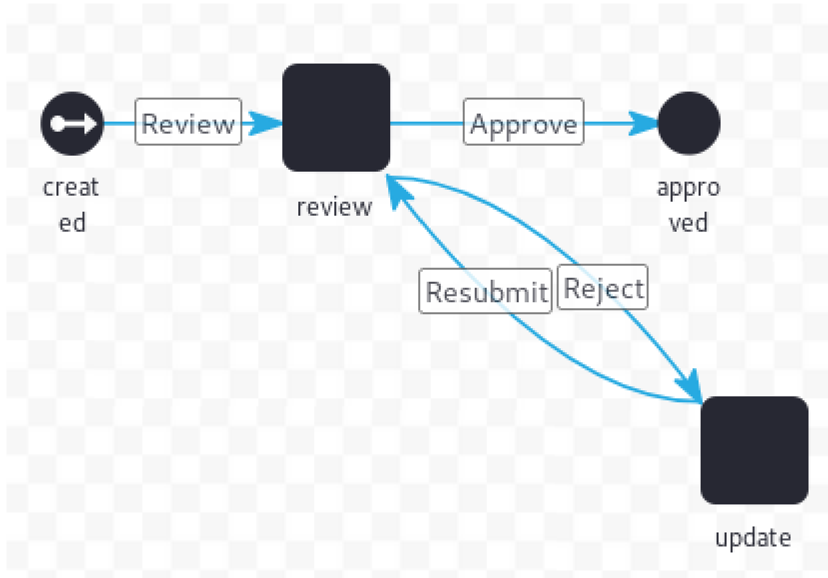

Figure 63.7: You connect nodes and direct workflow processing with transitions. The Single Approver workflow has transitions named Submit, Resubmit, Reject, and Approve.

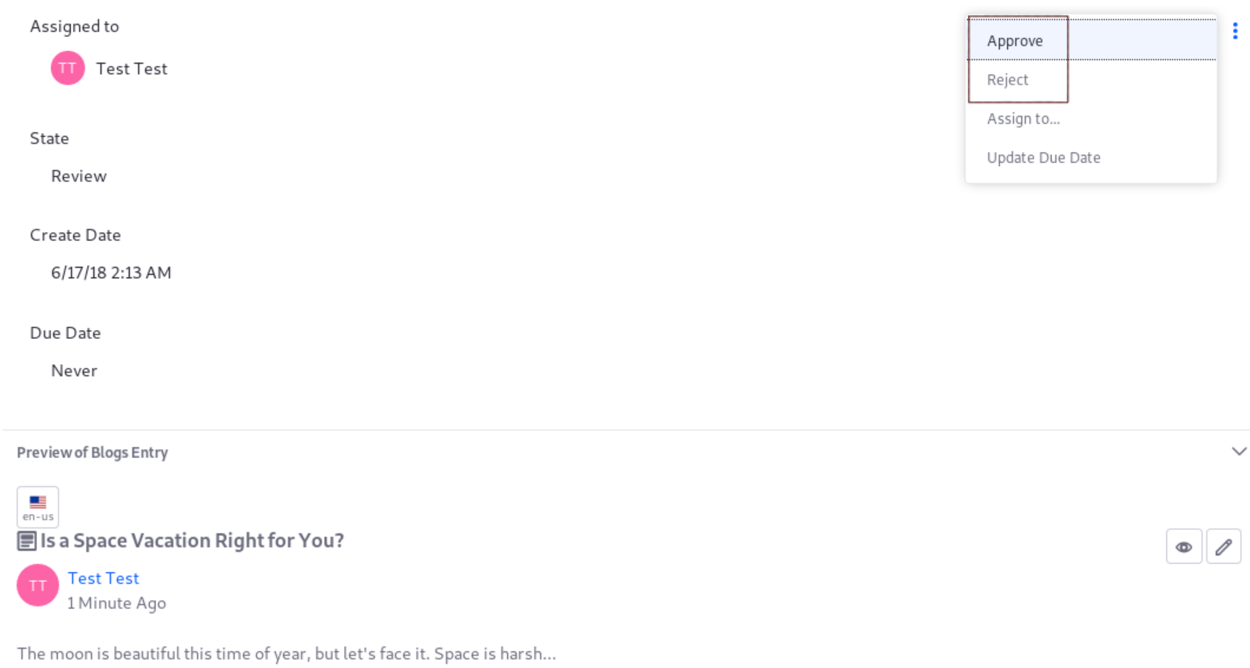

Figure 63.8: In the Single Approver workflow, a user in the Review task can choose to Approve or Reject the asset, which sends the asset either to the EndNode or to the Update task.

Forks and Joins

Sometimes you don't need to wait for one task to be completed before moving on to another one. Instead, you want to do two or more things at the same time. To do this, transition to a fork node, make two transitions from the fork to your parallel tasks, and then come back together using a join node.

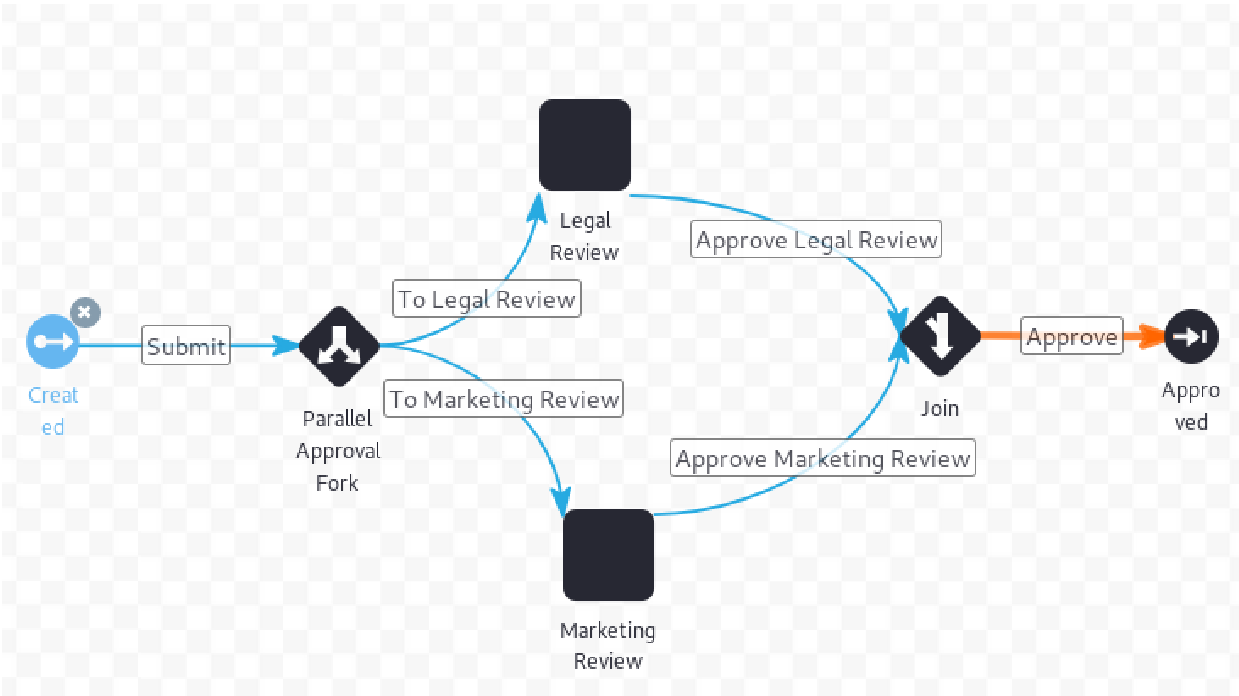

Figure 63.9: Forks and Joins are used to enable parallel processing in the workflow.

With a regular Join node, for the workflow to proceed beyond the join, the transition from both parallel executions must be invoked. However, if you use a Join XOR node instead, the workflow proceeds as long as the transition from one of the parallel executions is invoked.

Keep in mind that you must balance your fork and join nodes. In other words, for every fork, there must be a join that brings the parallel workflow threads back together.

Conditions

Sometimes you must inspect an asset or its execution context, and depending on the result, send it to the appropriate transition. You need a node for a script that concludes by setting a value to one of your transitions.

From the *Category Specific Approval* (`category-specific-definition.xml`), this is the script in the condition node that starts the workflow (coming directly from the start node):

```
import com.liferay.asset.kernel.model.AssetCategory;
import com.liferay.asset.kernel.model.AssetEntry;
import com.liferay.asset.kernel.model.AssetRenderer;
import com.liferay.asset.kernel.model.AssetRendererFactory;
import com.liferay.asset.kernel.service.AssetEntryLocalServiceUtil;
import com.liferay.portal.kernel.util.GetterUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;
```

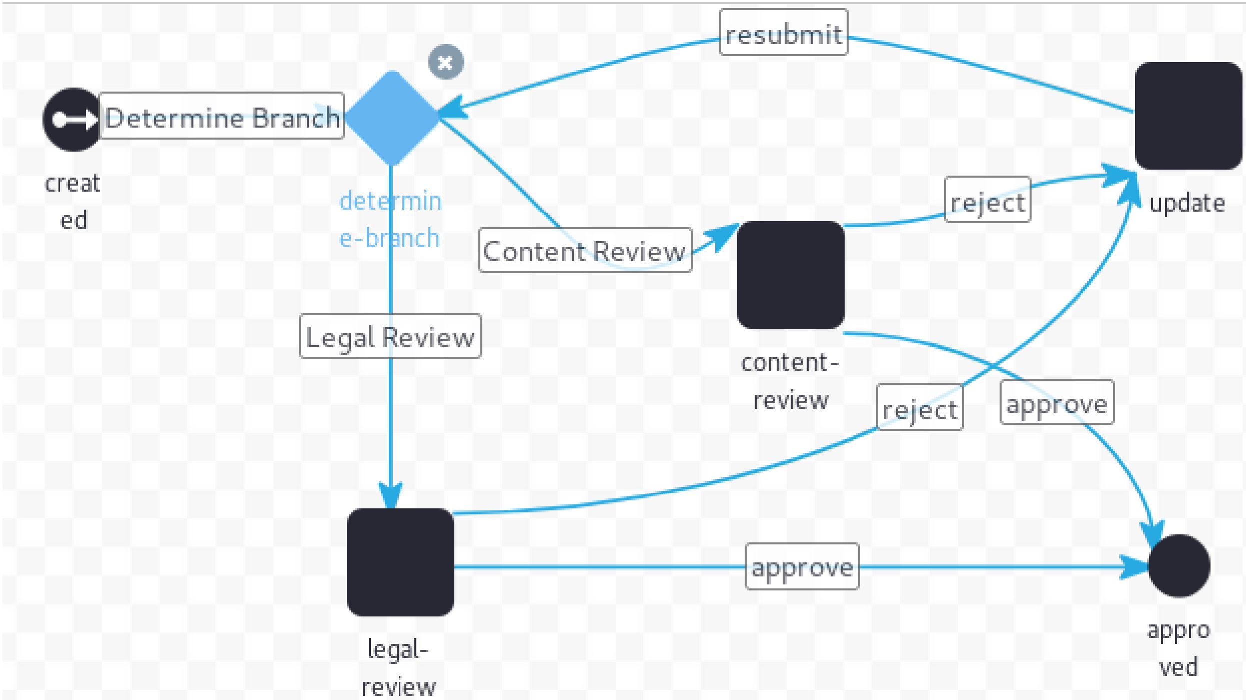

Figure 63.10: The Category Specific Approval definition starts with a Condition node.

```
import com.liferay.portal.kernel.workflow.WorkflowHandler;
import com.liferay.portal.kernel.workflow.WorkflowHandlerRegistryUtil;

import java.util.List;

String className = (String)workflowContext.get(WorkflowConstants.CONTEXT_ENTRY_CLASS_NAME);

WorkflowHandler workflowHandler = WorkflowHandlerRegistryUtil.getWorkflowHandler(className);

AssetRendererFactory assetRendererFactory = workflowHandler.getAssetRendererFactory();

long classPK = GetterUtil.getLong((String)workflowContext.get(WorkflowConstants.CONTEXT_ENTRY_CLASS_PK));

AssetRenderer assetRenderer = workflowHandler.getAssetRenderer(classPK);

AssetEntry assetEntry = assetRendererFactory.getAssetEntry(assetRendererFactory.getClassName(), assetRenderer.getClassPK());

List<AssetCategory> assetCategories = assetEntry.getCategories();

returnValue = "Content Review";

for (AssetCategory assetCategory : assetCategories) {
    String categoryName = assetCategory.getName();

    if (categoryName.equals("legal")) {
        returnValue = "Legal Review";

        return;
    }
}
```

This example checks the asset category to choose the processing path, whether to transition to the *Legal Review* task or the *Content Review* task.

You may be wondering what that `returnValue` variable is. It's the variable that points from the condition to a transition, and its value must match a valid transition in the workflow definition. This script looks up the asset in question, retrieves its asset category, and sets an initial `returnValue`. Then it checks to see if the asset has been marked with the *legal* category. If not it goes through *Content Review* (the content-review task in the workflow), and if it does it goes through *Legal Review* (the legal-review task in the workflow).

Now you're equipped with the basic knowledge to design beautiful, effective workflows so that your assets can be properly reviewed before they're published in your sites.

63.4 Creating Tasks in Kaleo Designer

Task nodes have several parts and are the most complex parts of a workflow definition. Unlike other workflow nodes, task nodes have Assignments, because a User is expected to *do something* (often approve or reject the submitted asset) when a workflow process enters the task node: the assignment specifies who that User is.

Commonly, task nodes contain Notifications, Assignments, and Actions (defined in scripts). See more about Notifications and Actions in the article on workflow nodes. Task nodes and their assignments are more complex and deserve their own article (this one).

To get started, drag and drop a task node on your workflow canvas if you haven't already. Open its Properties and give it a name. Then double click *Actions* in the task's Properties pane.

You can define a notification (often Task Assignee is appropriate), or write a Groovy script defining an action that's triggered for your task.

Next learn about creating Assignments for your task nodes.

Assignments

Workflow tasks must be completed by a User. You can choose how you want to configure your assignments.

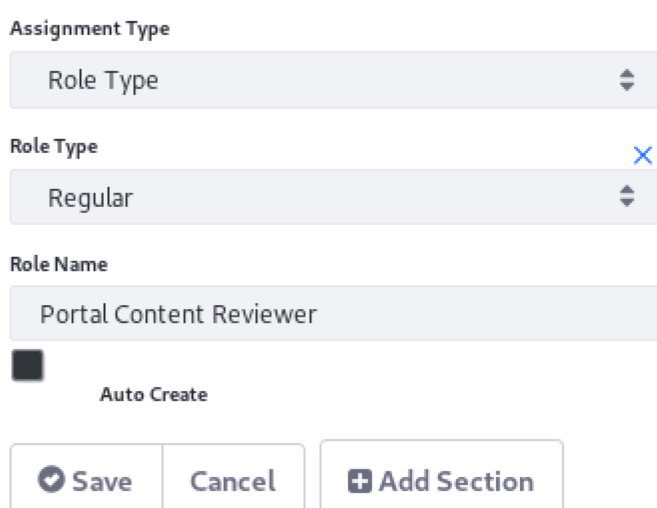

Figure 63.11: You can add an Assignment to a Task node.

You can choose to add assignments to specific Roles, multiple Roles of a Role Type (Organization, Site, or regular Role types), to the Asset Creator, to Resource Actions, or to specific Users. Additionally, you can write a script to define the assignment.

Assigning tasks to Roles, Organizations, or Asset Creators is a straightforward concept, but what does it mean to assign a workflow task to a Resource Action? Imagine an *UPDATE* resource action. If your workflow definition specifies the UPDATE action in an assignment, anyone who has permission to update the type of asset being processed in the workflow is assigned to the task. You can configure multiple assignments for a task.

Resource Action Assignments

Resource actions are operations performed by Users on an application or entity. For example, a User might have permission to update Message Boards Messages. This is called an UPDATE resource action, because the User can update the resource. If you're still uncertain about what resource actions are, refer to the developer tutorial on the permission system for a more detailed explanation.

To find all the resource actions that have been configured, you need access to the Roles Admin application in the Control Panel (in other words, you need permission for the VIEW action on the Roles resource).

- Navigate to Control Panel → Users → Roles.
- Add a new Regular Role. See the article on managing roles for more information.
- Once the Role is added, navigate to the Define Permissions interface for the Role.
- Find the resource whose action you want to use for defining your workflow assignment.

How do you go from finding the resource action to using it in the workflow? Use Kaleo Designer's interface for setting up a resource action assignment.

When configuring your task node's Assignment, select Resource Actions as the Assignment Type, then specify the Resource Actions to use for the assignment (for example, UPDATE).

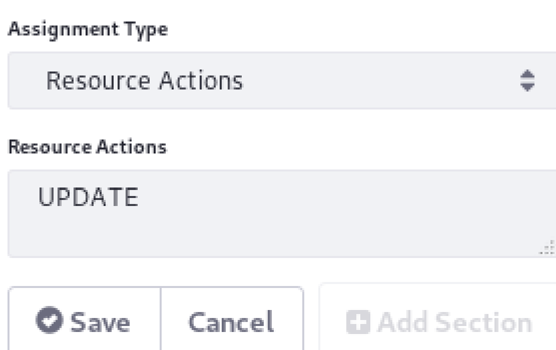

Figure 63.12: Configure resource action assignments in Kaleo Designer.

Here's what the assignment looks like in the Source (Workflow XML) tab:

```
<assignments>
    <resource-actions>
        <resource-action>UPDATE</resource-action>
    </resource-actions>
</assignments>
```

As usual, assign the workflow to the appropriate workflow enabled asset.

Now when the workflow proceeds to the task with the resource action assignment, Users with `UPDATE` permission on the resource (for example, Message Boards Messages) are notified of the task and can assign it to themselves (if the notification is set to Task Assignees). Specifically, Users see the tasks in their *My Workflow Tasks* application under the tab *Assigned to My Roles*.

Use all upper case letters for resource action names. Here are some common resource actions:

```
UPDATE
ADD
DELETE
VIEW
PERMISSIONS
SUBSCRIBE
ADD_DISCUSSION
```

You can determine the probable resource action name from the permissions screen for that resource. For example, in Message Boards, one of the permissions displayed on that screen is *Add Discussion*. Convert that to all uppercase and replace the space with an underscore, and you have the action name.

Scripted Assignments

You can also use a script to manage the assignment. Here's the script for the Review task assignment in the Scripted Single Approver workflow definition (`single-approver-definition-scripted-assignment.xml`):

```
import com.liferay.portal.kernel.model.Group;
import com.liferay.portal.kernel.model.Role;
import com.liferay.portal.kernel.service.GroupLocalServiceUtil;
import com.liferay.portal.kernel.service.RoleLocalServiceUtil;
import com.liferay.portal.kernel.util.GetterUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

long companyId = GetterUtil.getLong((String)workflowContext.get(WorkflowConstants.CONTEXT_COMPANY_ID));

long groupId = GetterUtil.getLong((String)workflowContext.get(WorkflowConstants.CONTEXT_GROUP_ID));

Group group = GroupLocalServiceUtil.getGroup(groupId);

roles = new ArrayList<Role>();

Role adminRole = RoleLocalServiceUtil.getRole(companyId, "Administrator");

roles.add(adminRole);

if (group.isOrganization()) {
    Role role = RoleLocalServiceUtil.getRole(companyId, "Organization Content Reviewer");

    roles.add(role);
}
else {
    Role role = RoleLocalServiceUtil.getRole(companyId, "Site Content Reviewer");

    roles.add(role);
}

user = null;
```

Don't let all that code intimidate you. It's just assigning the task to the *Administrator* Role, then checking whether the *group* of the asset is an Organization and assigning it to the *Organization Content Reviewer* Role if it is. If it's not, it's assigning the task to the *Site Content Reviewer* Role.

Note the `roles = new ArrayList<Role>();` line above. In a scripted assignment, the `roles` variable is where you specify any Roles the task is assigned to. For example, when `roles.add(adminRole);` is called, the Administrator role is added to the assignment.

63.5 Activating Workflow

Assets are integrated with the workflow framework if their content is meant to be sent through review processes. Since this is most often the case, there are many out of the box assets that integrate with workflow. In this article, you'll learn how to select a workflow for use with each of these workflow-enabled assets.

Workflow Assets

Activate a workflow for these assets in Control Panel → Configuration → Workflow, in the Schemes tab:

- Blogs Entry
- Calendar Event
- Comment
- Message Boards Message
- Page Revision
- User
- Web Content Article
- Wiki Page
- Knowledge Base

Activate workflow for these assets at the Site level in Site → Configuration → Workflow:

- Blogs Entry
- Calendar Event
- Comment
- Message Boards Message
- Page Revision
- Web Content Article
- Wiki Page
- Knowledge Base

What's the difference between the Site workflow configuration and the Control Panel Workflow configuration? As with most scoped configurations, the higher level setting (in the Control Panel) sets the default behavior. It's overridden at the more granular level (in the Site menu).

User doesn't appear on the Site list because adding users is strictly a portal-wide administrator activity. Only assets that can be added and configured at the Site level (for example, those that are accessed from the Site menu) have workflow configuration controls at the Site level.

Some assets that are workflow-enabled are activated in their respective application:

Activate workflow for Web Content Folders from the folder settings menu:

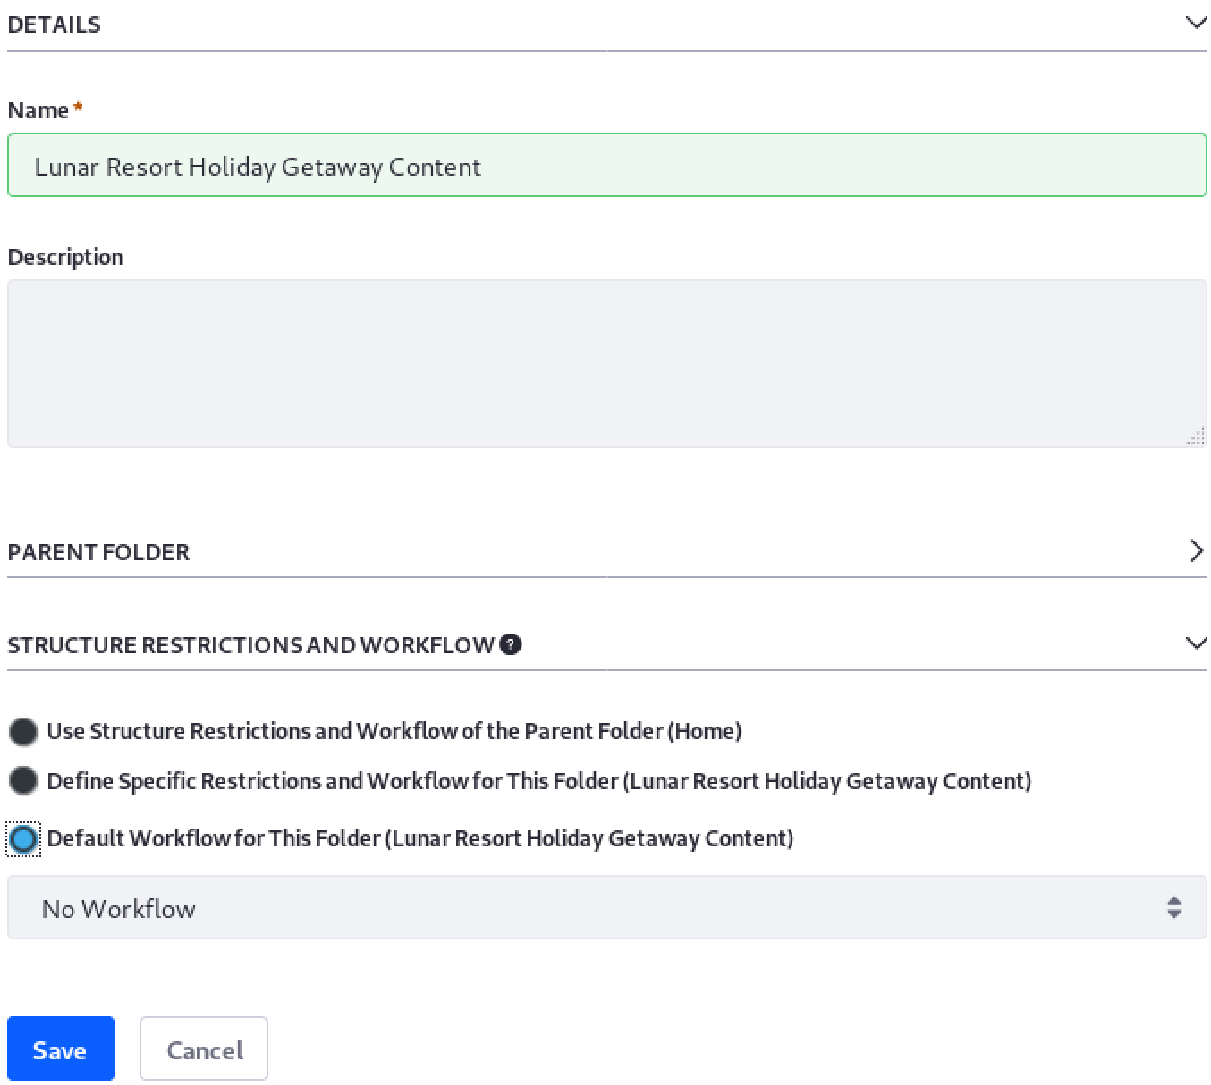

Figure 63.13: Activate workflow on Web Content folders from the folder's edit screen.

Activate workflow for Documents and Media Folders from the folder settings:

Enable workflow on Dynamic Data List entries in each list's Add form:

Activate workflow for each individual form's entries from the Form Settings screen:

Workflow Behavior

Most of the resources listed above behave just as you might expect with workflow activated: The Publish button for the resource's *Add* form is replaced by a *Submit for Publication* button, and instead of instant publication, the asset has its status set as *Pending* and must proceed through the workflow before publication.

Page revisions are slightly different. Page revisions only occur in staging environments that have Page Versioning enabled. When a Page Variation or Site Page Variation is created, its creator

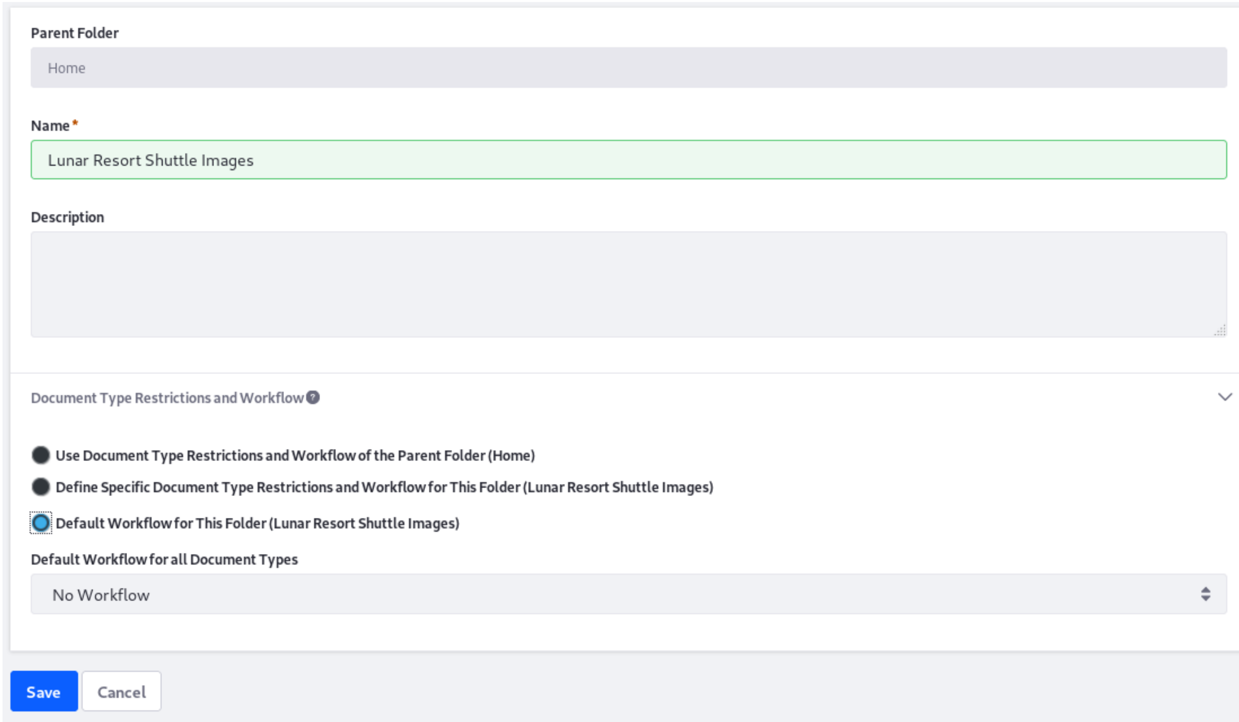

Figure 63.14: Activate workflow on Documents and Media folders from the folder's edit screen.

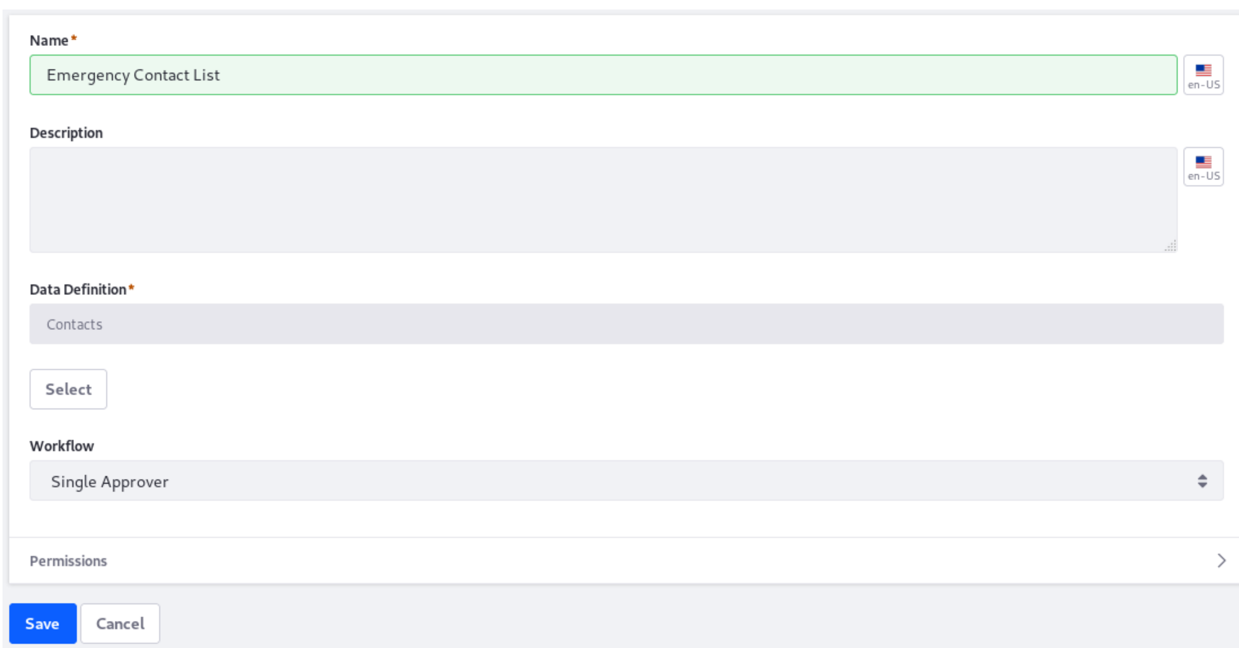

Figure 63.15: Activate workflow for each individual Dynamic Data List.

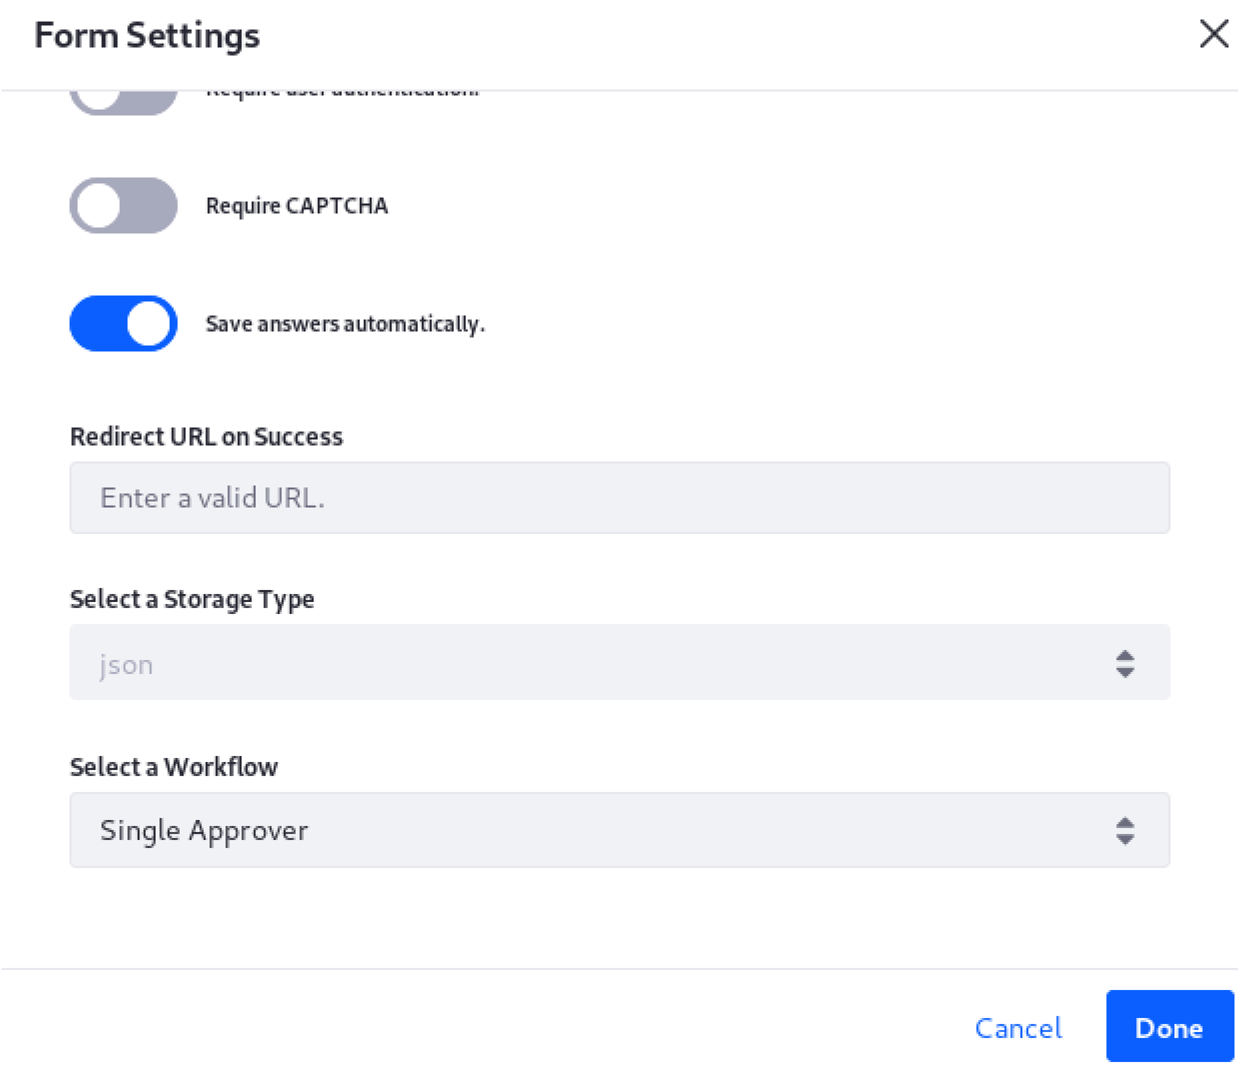

Figure 63.16: Activate workflow on each form's entries from the Form Settings window.

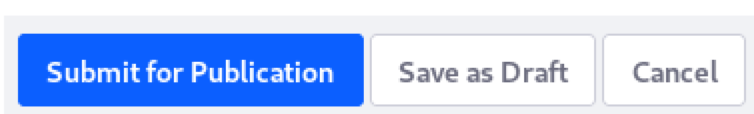

Figure 63.17: Instead of a Publish button, a Submit for Publication button appears for workflow-enabled resources.

must click *Submit for Publication* at the top of the page, and the variation must be approved in the workflow before it can be published to the live Site.

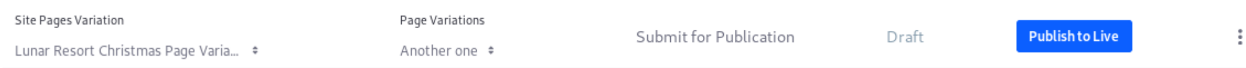

Figure 63.18: With workflow enabled on Page Revisions, the Site administrator must submit their page variation for publication before it can go live.

63.6 Kaleo Forms

Business processes are often form-based and workflow-driven. They start with entered data and progress by sending that data to other people or groups. Then it's processed in some way (for example, further data is entered or approval is granted), and the process moves on until completion, when each interested party has seen and manipulated the data. To write an app for each of these processes is laborious. It's much better to have a tool for quickly defining a process to suit each use case. The process architect must define both the data that gets collected and the process the data moves through to reach its final state. To accomplish this, Liferay DXP already includes the Dynamic Data Lists app for defining forms, and the Kaleo Designer app for designing workflows. The Kaleo Forms solution combines the features of these apps, letting you use a single UI to design an integrated process for sending forms through a workflow.

Creating Kaleo Forms Process

To start creating a Kaleo Forms Process you need to get to Kaleo Forms Admin:

1. Go to *Site Administration* (your Site's menu) → *Content* → *Kaleo Forms Admin*. The Kaleo Forms app appears with a list of any defined processes.

2. Click the *Add* button ( ) to open the New Process Wizard.

3. Give the process a name and a description and click *Next*.

4. Define the fields that can appear in your process's forms. There are two ways to do this:

   - Use an existing field set. Click the field set's Actions button ( ) and select *Choose*.
   - Create a new field set/data definition. Click the *Add Field Set* button. If you need help with this, see the documentation on creating data definitions. After creating the field set, select it as you would an existing field set.

   Click *Next* to move to the wizard's next step.

5. Select a workflow to use for your forms. To do this, click the workflow's *Actions* button ( ) and select *Choose*.

   You can also edit an existing workflow or create a new one:

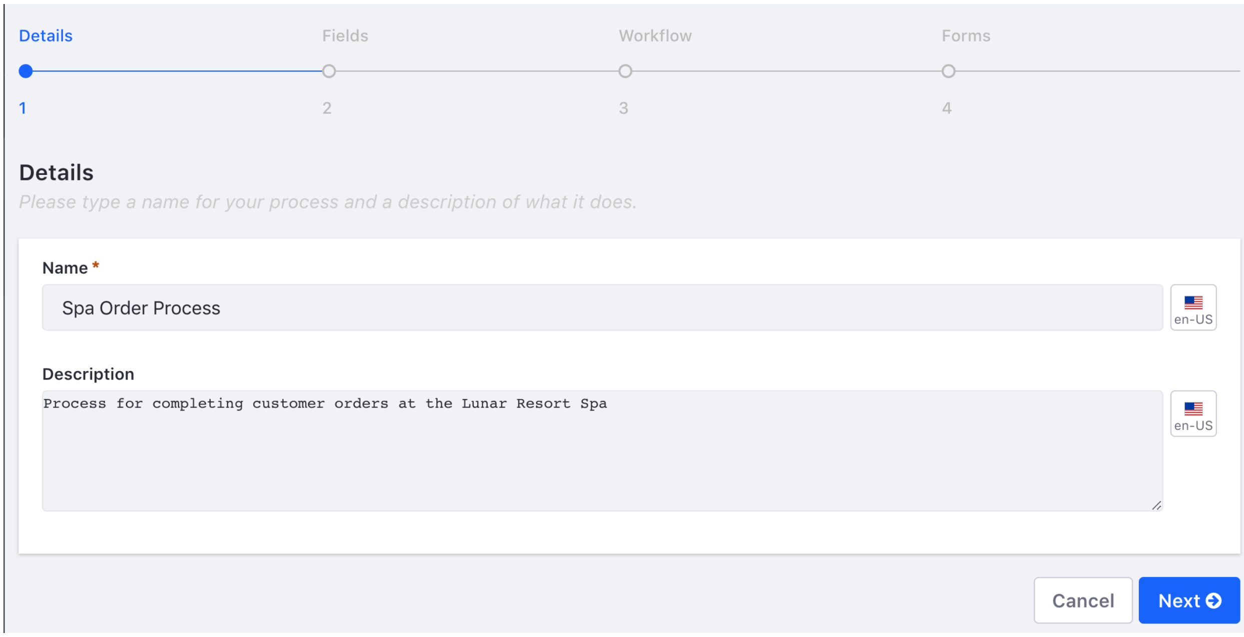

Figure 63.19: Add a Kaleo Forms Process to link a form with a workflow definition.

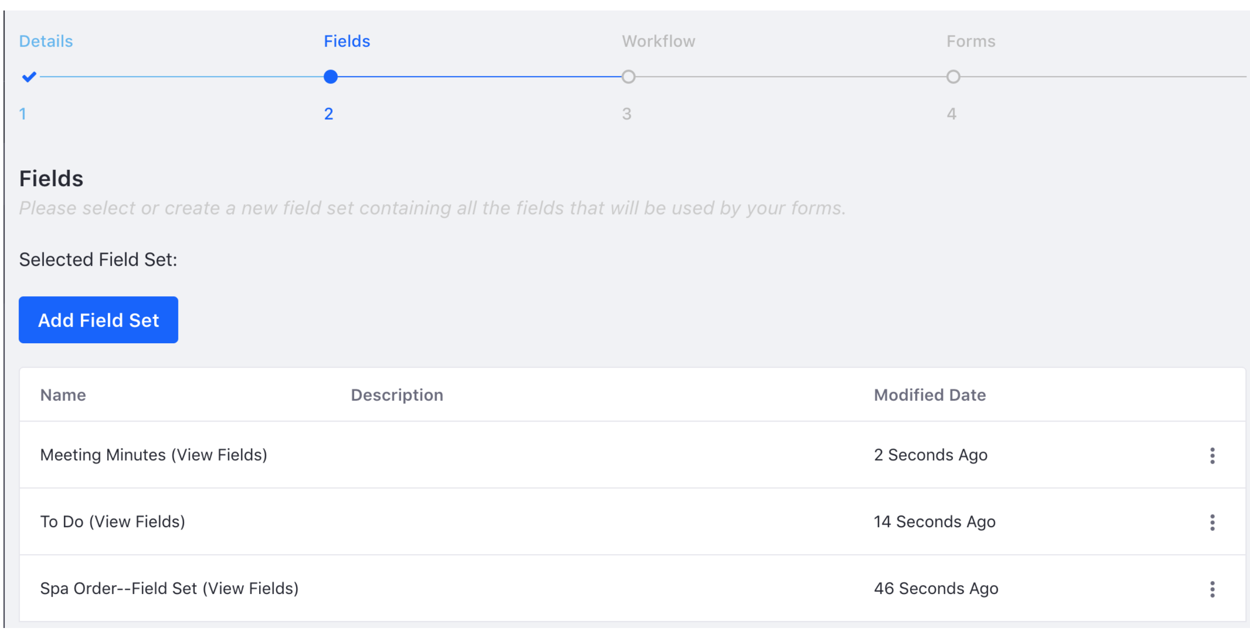

Figure 63.20: Define and choose your form's fields.

- To edit a workflow, click its *Actions* button (⋮) and select *Edit*.
- To begin creating a new workflow, click the *Add Workflow* button.

In either case, you use the same UI to edit/create the workflow. This UI is called Kaleo Designer. It lets you create your workflow graphically instead of via code.

Once you select a workflow to use with your forms, click *Next*.

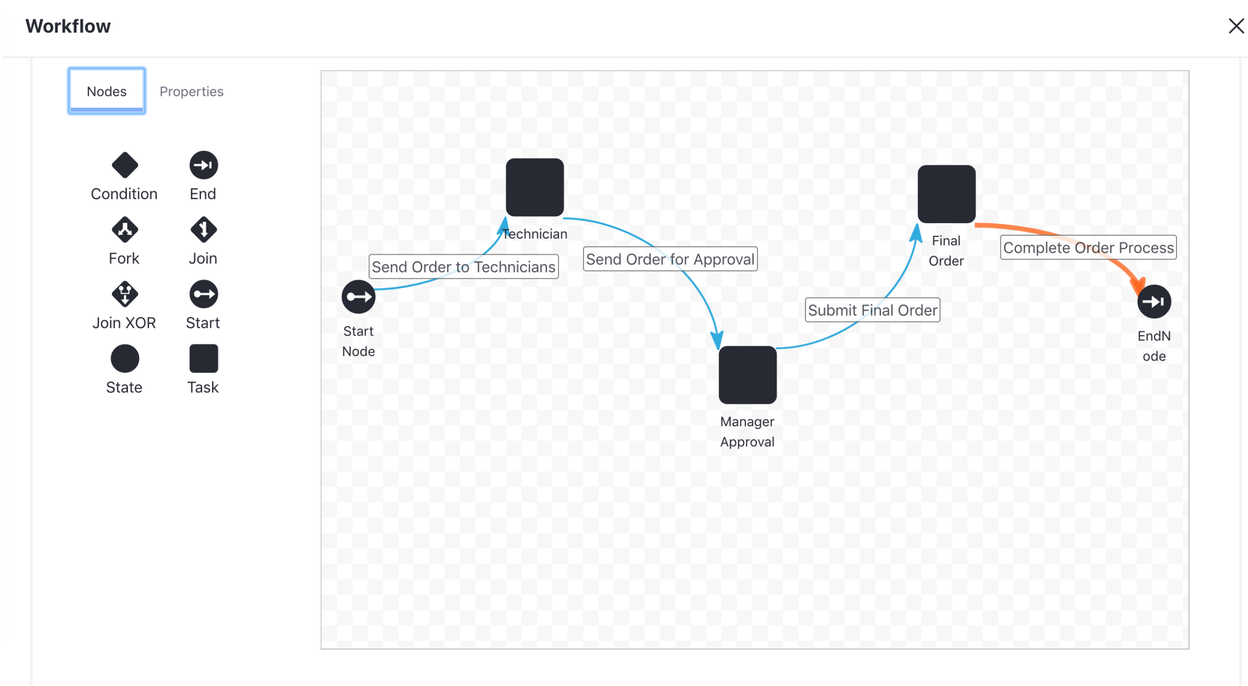

Figure 63.21: This example workflow has three tasks that happen sequentially.

6. Select or create a form to use for each workflow task. To do this, click each task's *Actions* button (⋮) and select *Assign Form*. On the screen that appears, select an existing form or click the *Add* button (+) and to create one.

   Click *Save* when you're finished. Your process is done and appears in Kaleo Forms Admin's table.

Adding Records to a Process

To add records to a process, click it in Kaleo Forms Admin and then click the *Add* button (+). This brings up the form you assigned to the workflow's initial state. Fill it out and click *Save*. Once submitting the initial form, the workflow engine then takes over and moves through each task in the workflow. Whatever Users or Roles you assigned to the tasks receive notifications, and the task appears in the *Assigned to My Roles* section of the My Workflow Tasks app. A notification also appears in the Notifications app. Once in the task, the user views and approves the form or clicks the *Edit* button. At this point, the workflow task forms you created come into play. Each assigned user fills out the form, saves it, and sends it along in the workflow.

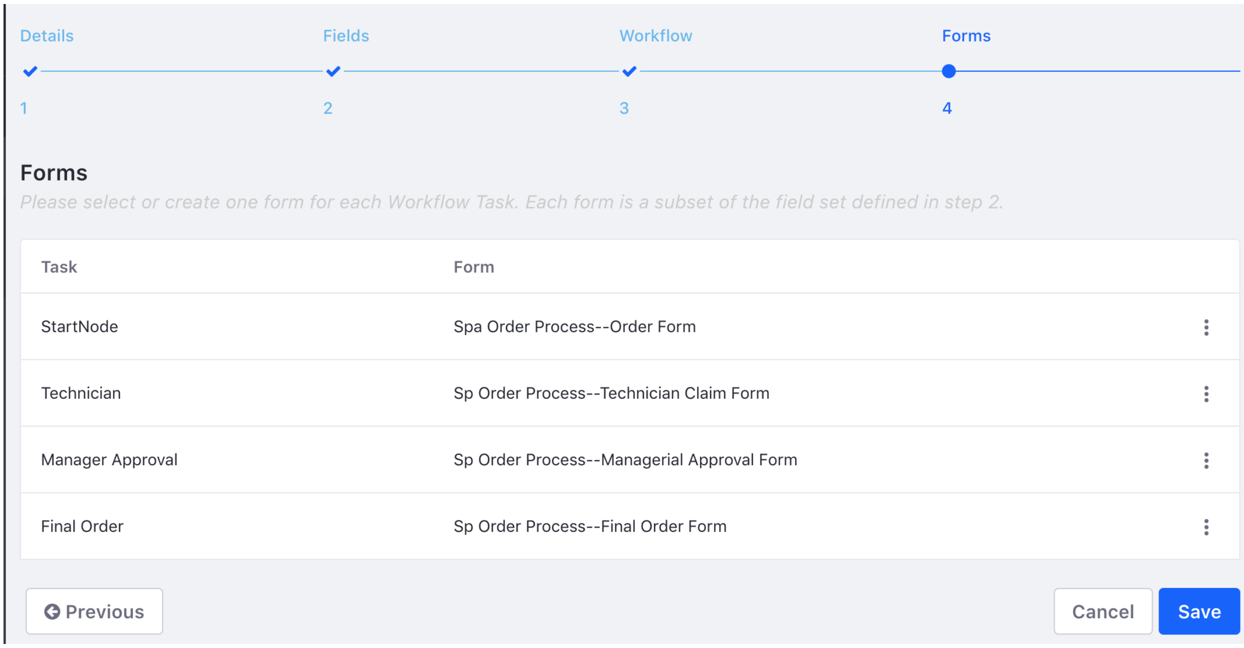

Figure 63.22: Assign a form to each task in the workflow, and for the initial state.

63.7 Managing Workflows

Like other resources, workflow definitions can be added, edited, and deleted. But that's just the beginning of workflow management.

- Because workflow definitions can be complex works in progress, they can be versioned.
- Unpublished drafts can be saved.
- Because workflow definitions are XML files, they're portable. Thus, they can also be uploaded.

Start by learning about the permissions for editing and publishing workflow definitions.

Workflow Definition Publication Permissions

Users with permission to edit or publish workflow definitions can add Groovy scripts to the workflow. Access to the scripting engine means access to the Java Virtual Machine (JVM) of the server. Users who publish (or edit) workflow definitions containing scripts, therefore, can get access to any data within the reach of the JVM, such as data contained in a separate Virtual Instance of Liferay DXP itself.

Because of this far-reaching access, permission to create or edit workflow definitions is limited to Regular Administrators of the Default Virtual Instance. For Liferay DXP customers, Site Administrators who can access the Kaleo Designer and Kaleo Forms Admin applications can also edit and publish workflows.

To grant Users with these Roles the workflow publication access in additional Virtual Instances,

1. Make sure you understand the access you're granting these admins.
2. Navigate to Control Panel → System Settings → Workflow → Workflow Definition.
3. Check the box for the setting *Allow Administrators to Publish and Edit Workflows*.

This only applies to Virtual Instances that have been added to the system. The Default Virtual Instance provides workflow publication access to Regular Administrators, via Control Panel → Configuration → Workflow and in Liferay DXP with access to the Kaleo Designer and Kaleo Forms Admin applications.

Adding, Editing, and Deleting

To add a workflow definition,

1. Go to Control Panel → Configuration → Workflow.
2. In the default view (Workflows), click the Add button (+).
3. From here you're either writing an XML definition, designing a definition in Kaleo Designer (DXP only), or uploading an existing definition.

To edit a definition,

1. Go to Control Panel → Configuration → Workflow.
2. Click the *Actions* button (⋮) for the workflow, and click *Edit*.

To delete a definition,

1. Go to Control Panel → Configuration → Workflow.
2. Click the *Actions* button (⋮) for the workflow, and click *Unpublish*.

   A published workflow cannot be deleted, so you must unpublish its workflow definition first.

   You can't unpublish a definition if it's activated for an asset. First dissociate the workflow definition from any assets that use it. See here for more information.

Uploading Workflow Definitions

If you have a local XML definition file (perhaps you want to create a new workflow based on one of the embedded workflows), upload it to Liferay DXP:

1. Navigate to Control Panel → Configuration → Workflow.
2. Click the *Add* button (+).
3. Name the workflow; then click the link to *import a file*.
4. Find the XML file and upload it. Once uploaded, the definition's XML appears in the workflow editor.
5. If it's ready to publish, click *Publish*. Otherwise, *Save* it and it stays Unpublished.

Liferay DXP customers use the Kaleo Designer to create workflows by default. To upload existing XML files in Kaleo Designer,

1. Navigate to Control Panel → Configuration → Workflow.
2. Click the *Add* button (+).

   The Kaleo Designer palette appears.
3. Click the *Source* tab.
4. Click the hyperlink *import a file* in the sentence

   `Write your definition or import a file`
5. Find the XML file and upload it. Once uploaded, the definition's XML appears in the workflow editor.
6. To edit using the Kaleo Designer drag-and-drop interface, click the *Diagram* tab.
7. If it's ready to publish, click *Publish*. Otherwise, *Save* it and it stays Unpublished.

What's the difference between saving and publishing?

Published Versus Unpublished

The difference between a published and unpublished workflow is important:

Published: Validation is complete, and the workflow can be assigned to assets.

Unpublished: Validation is not performed on the unpublished workflow, and it cannot be assigned to assets until it's published.

Workflow Versions

You're making a simple edit to a workflow, when suddenly you remember you have a meeting with your boss. Quickly you save the workflow and hurry off to your meeting. Congratulations! You were promoted to Director of Business Productivity! You have no time to edit workflows now, so your colleague must finish editing and publishing the workflow. Unfortunately, in all the excitement of your promotion, you forgot what you changed in the workflow. It's best to revert to the prior version and start editing it from scratch.

1. Open the workflow editor. Go to *Control Panel* → *Configuration* → *Workflow*, and select the workflow from the list.
2. Open the information panel. Click the *Information* button ( )
3. There are two information panel sections: Details and Revision History.

   The Details screen shows information about the creation of the workflow and its last modification, and a summary of the total modifications.

   The Revision History screen shows the workflow's current and prior, restorable versions. To view an old workflow or to restore it if you're sure it's the right version, click the *Actions* button ( ) and select either *Preview* or *Restore*.

4. When you click *Restore* and see the success message, the prior version is the current version of the workflow. You can now edit the restored version of the workflow.

5. If edits are necessary, edit and click *Update*. This creates another version of the workflow.

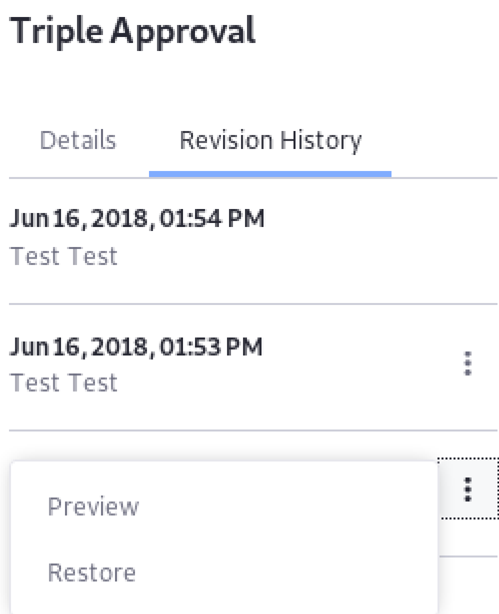

Figure 63.23: View and restore prior versions of a workflow.

Alternatively, you can refer to the embedded definitions to get workflow definition ideas.

63.8 Reviewing Assets

User interaction is required at each workflow process Task Node. How do users complete tasks? In the *My Workflow Tasks* application of the *My Account* menu.

Asset Submission to Workflow

If an Asset has a workflow activated, when a user clicks Submit for Publication, the workflow definition determines the next step. A user assigned a Role associated to the workflow task receives a Notification indicating that there is a workflow task to complete.

Workflow Assignees Have Implicit Asset Permissions: Users with permission to execute a workflow task (e.g., Users with the Portal Content Reviewer Role) have full resource action permissions over the assets they can review. These permissions apply in the My Workflow Tasks widget in the User's personal page and anywhere else actions on the Asset can be performed.

For example, consider a User with two permissions:

- The Portal Content Reviewer Role enables Users to review workflow submissions and grants edit and delete permissions on the content they're reviewing.
- Users also have permission to view Web Content Articles in the Site's *Content* section.

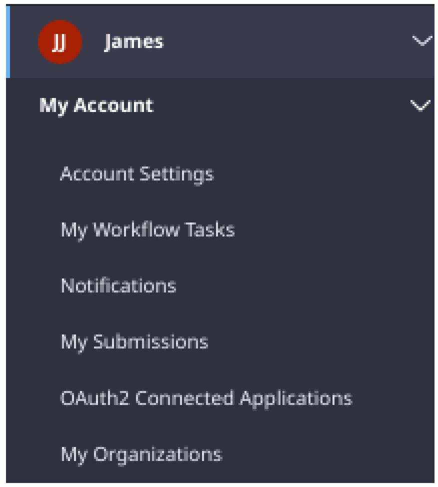

Figure 63.24: Users manage workflow tasks from their My Workflow Tasks widget.

Neither permission explicitly grants the User management permissions on Web Content Articles. Users cannot normally edit or delete a Web Content Article, for example. However, if a Web Content Article is sent to the workflow, Users can access the Web Content Article for review (in their *Assigned to Me* or *Assigned to my Roles* section of My Workflow Tasks), and they can edit or delete the content while reviewing it in the workflow. While it's in the status *Pending*, they can also edit or delete the article from Site Administration → Content → Web Content because of their implicit permissions granted by the workflow system. This additional permission is temporary, and the normal resource permissions are activated once the Web Content Article exits the workflow process (for example, it's rejected or approved).

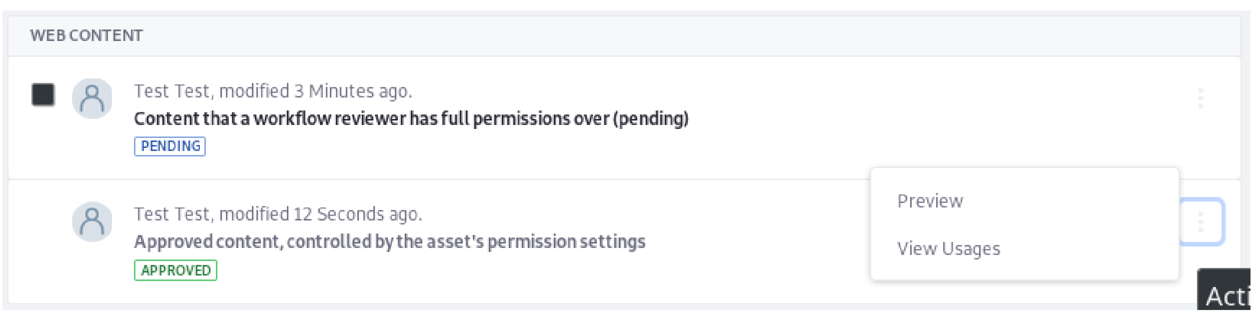

Figure 63.25: A User with VIEW permission on Web Content cannot manage Approved Articles.

Assigning the Task

Workflow Tasks can be completed only by certain users, based on the Assignment.

All workflow tasks assigned directly to a user are listed in the My Workflow Task widget's *Assigned to Me* tab.

If a workflow was assigned to a Role that the user occupies, the workflow's tasks appear in the *Assigned to My Roles* tab.

To claim a task, the user must move the task into the *Assigned to Me* tab.

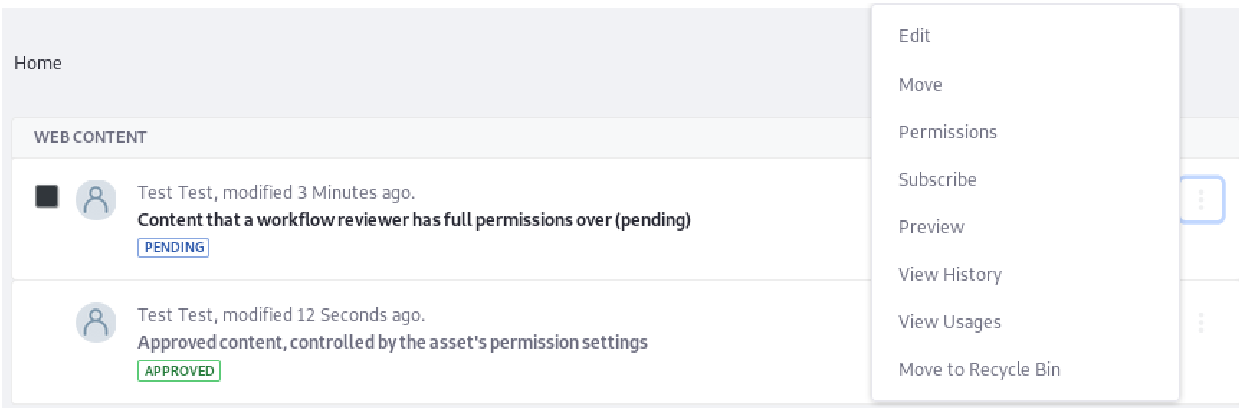

Figure 63.26: A User with access to Web Content in the Workflow can manage Pending Articles.

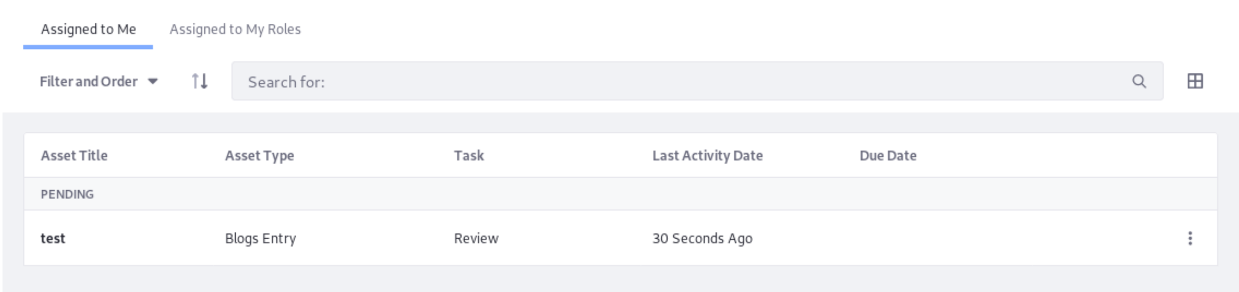

Figure 63.27: The assets assigned to a user are listed in *Assigned to Me*.

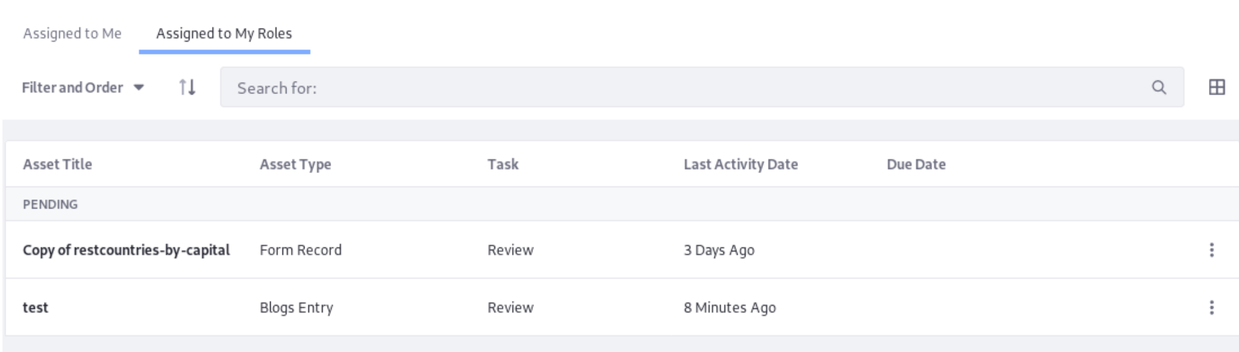

Figure 63.28: The Assets assigned to Roles are listed in each associated user's *Assigned to My Roles* tab.

1. Click the Asset's *Actions* button ( ) and select *Assign to Me*.
2. Add a comment in the pop-up box if necessary, and click *Done*.

Alternatively, assign the task to another user.

1. Click the Asset's *Actions* button ( ) and select *Assign to ...*.
2. Select the user to assign the task.
3. Add a comment in the pop-up box if necessary, and click *Done*.

Completing the Task

Once a task is assigned, it's ready to be completed.

There's a fast way to send an asset along in the review process:

1. From the *Assigned to Me* tab, click the *Actions* button ( ) and select *Approve* or *Reject*.

   Note that options names in this menu are identical to the the workflow definition's Transition names. Your menu might have different options than the *Approve* and *Reject* options in the figure below.

2. Enter a comment if desired and click *Done*.

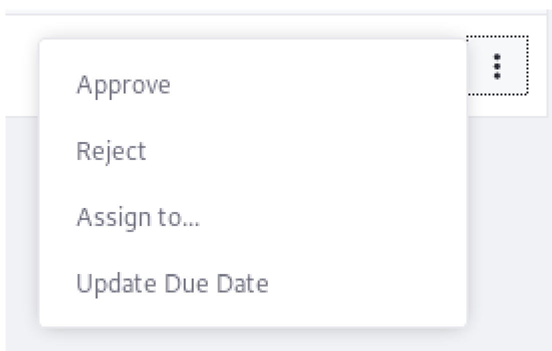

Figure 63.29: Complete Tasks right from the *Assigned to Me* list.

Here's how to get a closer look at the Asset before sending it along in the workflow:

1. From the *Assigned to Me* tab, click the title of the Asset to review.

   The Task screen appears showing details about the Asset:

2. Inspect the Asset to your liking (or even edit it if you have permission) and click the *Actions* button ( ).
3. Send it along in the workflow by clicking one of the Transition names (for example, *Approve* or *Reject* in the Single Approver Definition).

And, you're done. Once you've completed your workflow tasks, kick back and wait for more to come in.

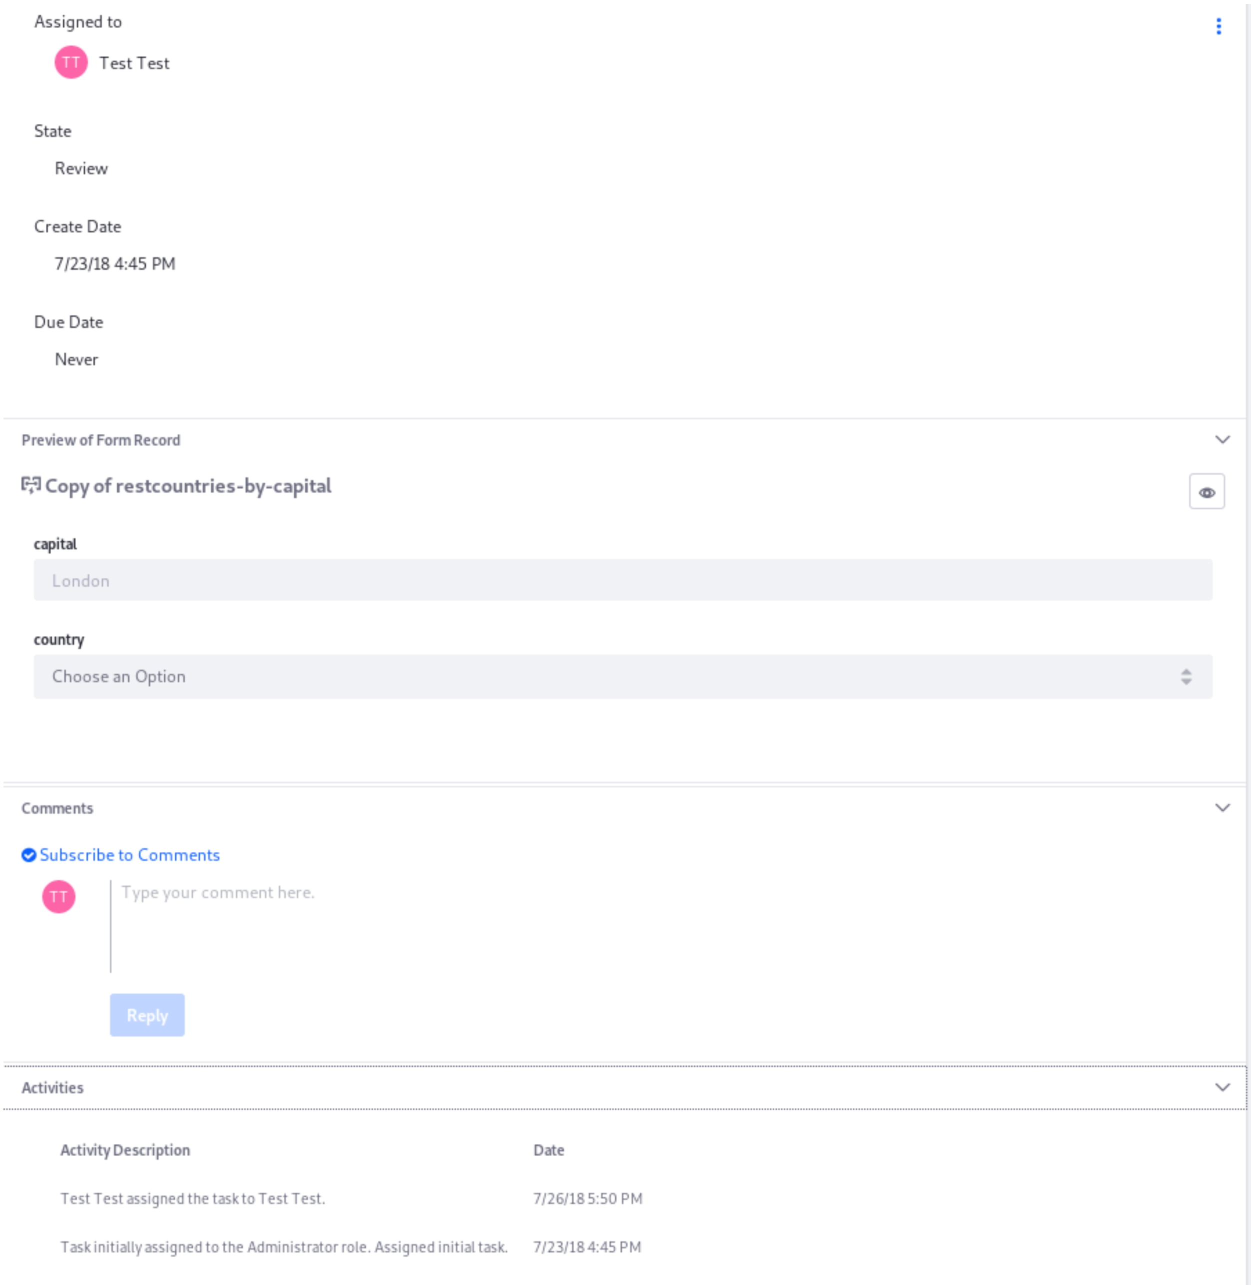

Figure 63.30: Inspect Assets before completing the Task.

Chapter 64

Targeting Content To Your Audience

Audience Targeting provides tools to customize user experience based on demographics, browsing data, or other criteria. You can divide your audience into user segments, target content to them, and create campaigns to get content to the right people.

With Liferay's Audience Targeting app, developers can easily create additional rules and rule types with minimal coding effort. Campaigns last for fixed periods of time and each campaign has a priority that determines which campaign takes precedence, and you can configure different assets to be displayed at different periods of time to the targeted user segment.

You can also measure the effectiveness of your campaigns. Different assets are shown to different users, and campaigns display specific content to different user segments for fixed periods of time, measuring the interaction of the targeted user segments with the chosen content.

You can use Audience Targeting to display different sets of advertisements to different user segments, and this can dynamically change every week based on your marketing goals. Or you can display news articles to users based on all the user segments to which they belong.

Audience Targeting meets all of the requirements for anonymization, erasure, and export for compliance with GDPR.

The Audience Targeting App appears in the Configuration section of Site Administration. It also adds a User Segments Simulator in the Simulation Menu. The following three applications are also included with the Audience Targeting app:

- Campaign Content Display
- User Segment Content Display
- User Segment Content List

You can add these widgets to any Liferay DXP page. In the next sections, you'll lean how to use the *Audience Targeting* section of the Site Administration area to manage user segments and campaigns. Then you'll learn how to use each of the Audience Targeting applications and the Audience Targeting Simulator.

Note: The IP Geocoder porlet is an Optional Dependency of Audience Targeting and while it is not required to deploy Audience Targeting, you will need it to access the full functionality of Audience Targeting. For more information on Optional Dependencies, see our Declaring Optional Import Package Requirements troubleshooting page.

64.1 Managing User Segments

Go to *Site Administration* → *Configuration* → *Audience Targeting* to manage user segments and campaigns.

A user segment represents a subset of the total group of users (logged in or not). A user segment is defined by one or more rules that users must match to belong to that user segment. Once the user segment is created, only users who visit the applicable Site(s) are added to it. To add a new User Segment,

1. Open Site Administration.
2. Go to *Configuration* → *Audience Targeting*.
3. Under the *User Segments* tab, click *Add User Segment* ().
4. Enter "Women over 30 who use Safari" in the *Name* field and enter a *Description*.
5. Open *Behavior*. Click and drag *Browser* to the right side.

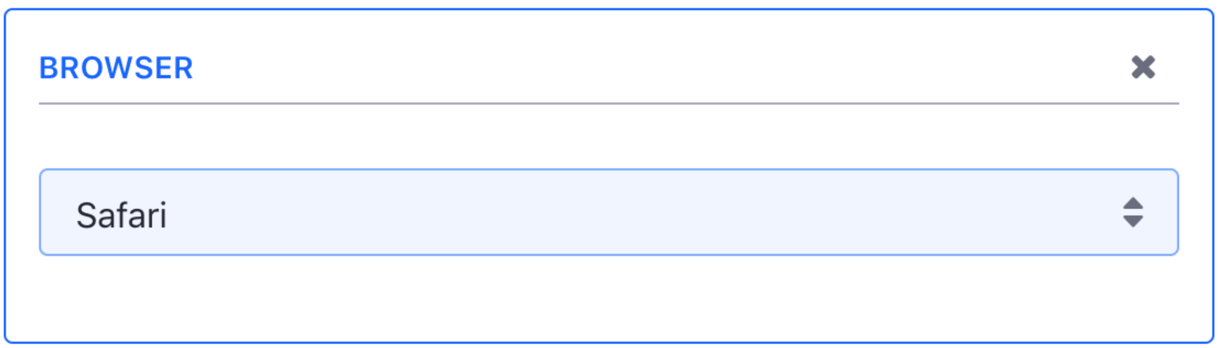

Figure 64.1: After applying the rule, click the rule to expand/collapse it.

6. Select *Safari* in the *Browser* selector.
7. Open *User Attributes* and drag *Age* and *Gender* under where you set the browser.
8. Set the *Age* to older than 30 and younger than 100.
9. Set the *Gender* to *Female*.
10. Click *Save*.

Each rule has parameters that can be adjusted. Those parameters can be in different formats. Some rules are also instanceable, meaning you can apply more than one of the same type of rule to a user segment for scenarios where an audience must be segmented based on actions that might occur multiple times (e.g., visiting multiple pages, viewing multiple banners, having several roles, etc.).

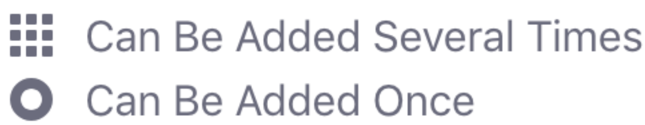

Figure 64.2: Instanceable and non-instanceable rules have different icons.

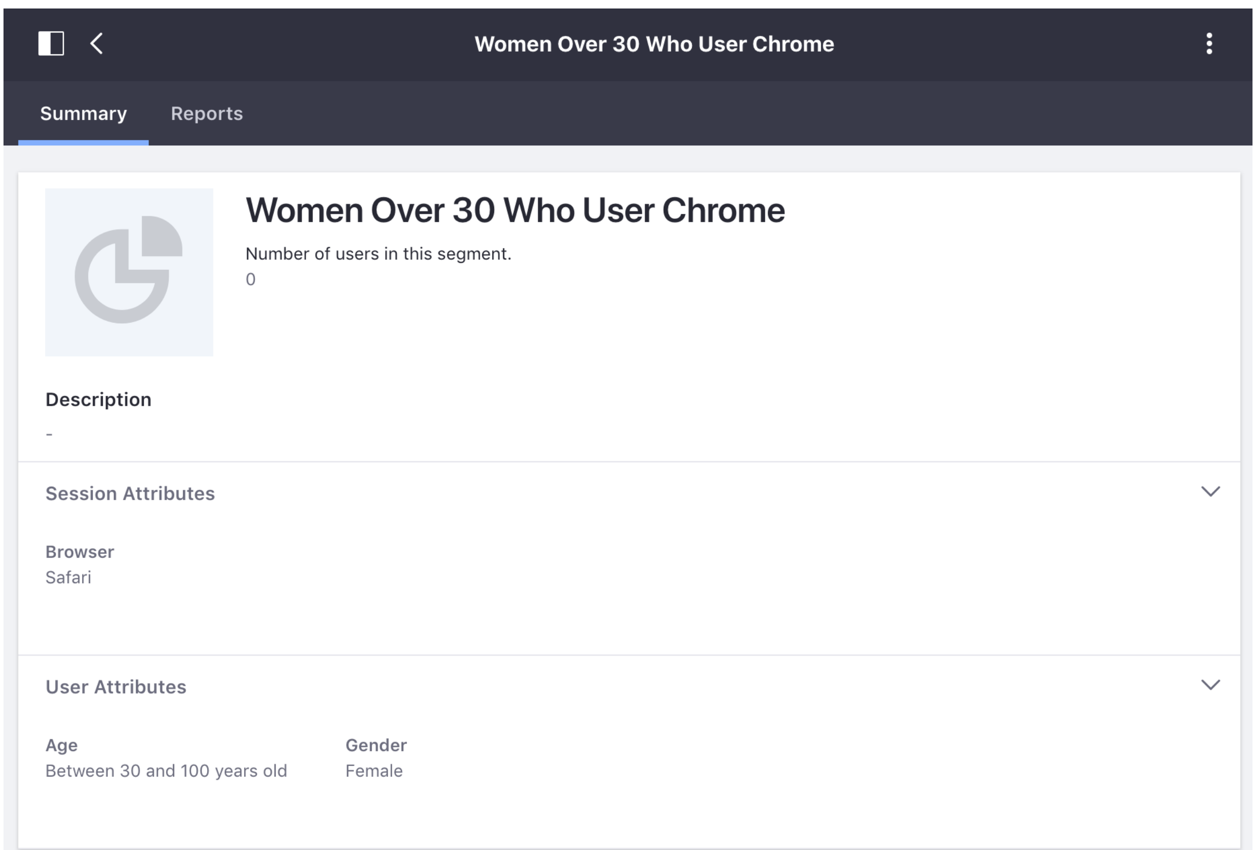

Figure 64.3: Select a pre-existing user segment to view its Summary page.

Once you've created a user segment, you can open its summary view, which displays relevant data and configurations.

To configure permissions for creating, editing, and deleting user segments,

1. Go to the *Users → Roles* section of the Control Panel.
2. Click on the *Actions → Define Permissions* button corresponding to the Role you want to configure permissions for, and search for *Audience Targeting*.

Both user segments and campaigns are inherited in your Site hierarchy. Therefore, user segments and campaigns defined in the Global scope are available to any Site. User segments and campaigns created in a site are available to all of its child Sites.

User Segment Rules

Custom rules can be created by developers and deployed as OSGi plugins. See the Creating New Audience Targeting Rule Types tutorial for details.

These are some of the rules that are included with the app by default:

- User Attributes
  - Age (from the user profile)
  - Gender (from the user profile)
  - Role (regular Role, Organization Role or Site Role)
  - Membership (Site member, Organization member, User Group member)
- Session Attributes
  - Location (obtained from the IP address)
  - Browser, Device, Operating system
  - Time
- Social
  - Like of a specific Facebook page
  - Number of Facebook friends
  - City, Age, Gender, Education, etc. from your Facebook profile
- Behavior
  - Viewed page or content
  - Score Points rule.

The Score Points rule assigns one point to a user each time the user views a page or content categorized under the user segment where the rule applies. Once a user exceeds the configured threshold, the user matches this rule.

For example, suppose that your website is about sports and you have pages and content about basketball, tennis, and soccer. You want to divide your audience into three user segments (Basketball fans, Tennis fans, and Soccer fans) to display the most relevant content to them on your site's front page. After creating these three user segments using the Score Points rule with a threshold of, say, 20, you should appropriately categorize the content most relevant to each user segment. For example, apply the *Basketball fans* user segment to content about basketball, apply the *Tennis fans* user segment to content about tennis, etc. Now every time a user (even a guest user) visits a page or views a piece of content categorized for a user segment where the Score Points rule applies, the user starts accumulating points. Once the user has accumulated enough points, the user belongs to the user segment. After a user has visited more than 20 pages or pieces of content related to basketball, the user belongs to the Basketball fans user segment. Once the user belongs to a user segment, you can use that information to direct more relevant information to the user using the User Segment Content Display application.

64.2 Categorizing Pages and Content for User Segments

Each new user segment that's created can be used to categorize pages or content. The Audience Targeting app adds a new *User Segment* select button to the SEO section of pages and Metadata section for assets. These buttons can assign one or more Site-scoped or global user segments to the content. This categorization has mainly two purposes:

- Assigning points to users using the Score Points rule
- Showing dynamic lists of content in the User Segment Content List application

You don't have to create categories for each of your user segments. User segments are distinct from regular vocabularies. The editing screen for both pages and assets contain distinct select buttons for user segments and regular vocabularies.

Another way to display user segments is through the Asset Publisher app. You can enable the Asset Publisher to retrieve assets that have matching categorization with the user segments of the current user. This enhances the Asset Publisher to only display relevant content to the user. To enable *User Segments Filter*,

1. Navigate to the Asset Publisher's *Options* ( ) → *Configuration* menu.
2. Under the *Asset Selection* tab, Open the *User Segments Filter* option.
3. Enable the *User Segments Filter*.

Next, you'll learn about managing user segment reports.

Managing User Segment Reports

When managing user segments, you can select the user segment name and then select the *Reports* tab to see the list of reports available for each user segment. Click the report name to view the report or *Actions* ( ) → *Update Report* to generate a new report. Reports display a summary of interesting information related to each user segment. For example, the Content Views report shows the asset viewed the most by users that belong to the user segment.

Reports also display which users belonged to a user segment. This lets an administrator know which users of the Site joined the particular user segment. You can export this list by clicking the user report you're interested in and selecting the *Options* () → *Export* button. This downloads a CSV file with the list of users.

Additional reports can be created by developers and deployed as OSGi plugins. Reports are generated daily by default, but you can generate a report at any time. To generate a new report

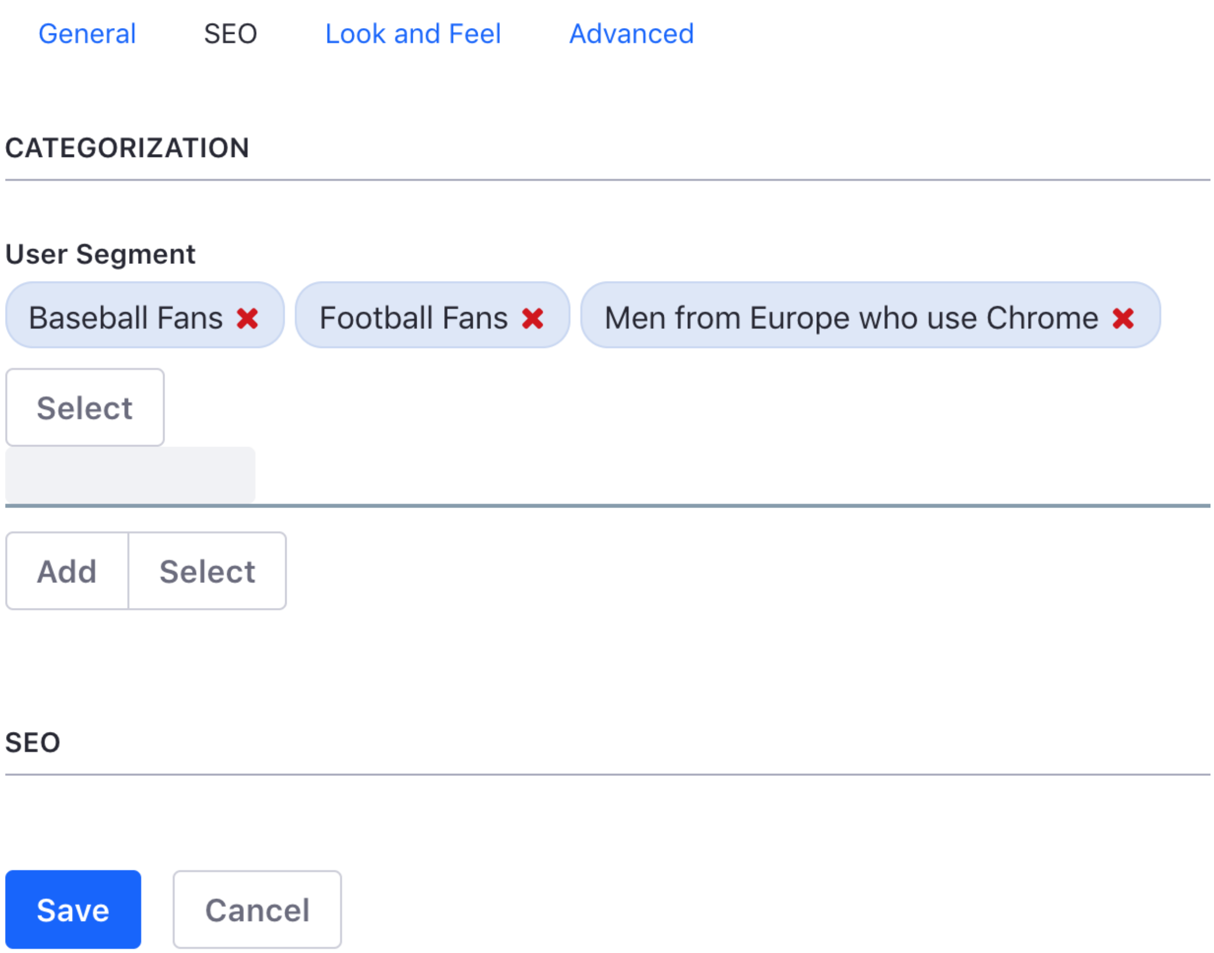

Figure 64.4: Pages and content can be categorized for user segments.

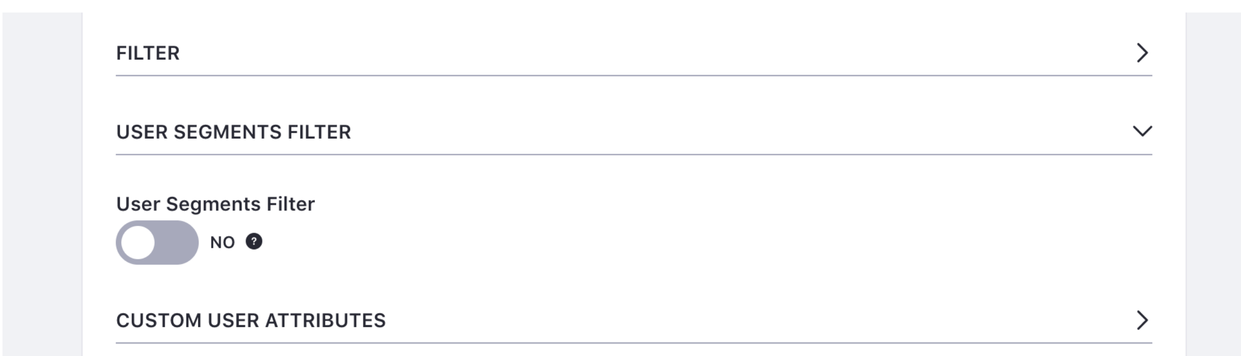

Figure 64.5: Enabling the User Segments Filter retrieves assets that match the current user's user segments.

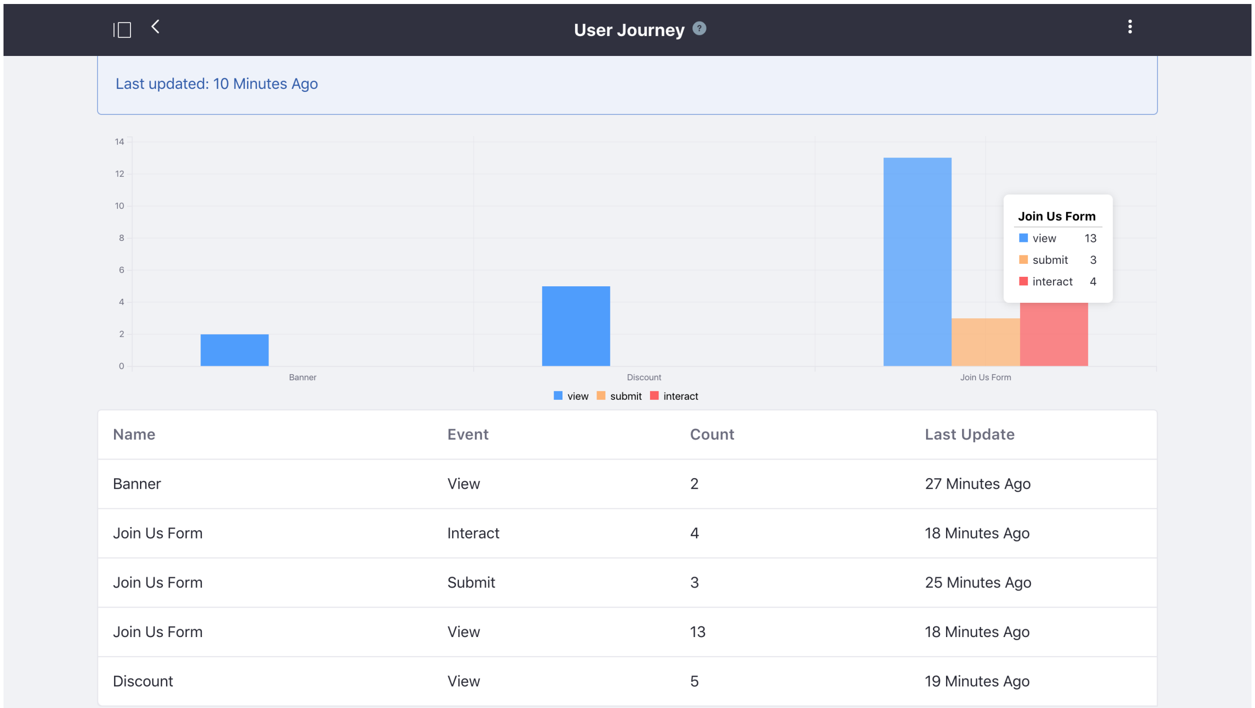

Figure 64.6: This report displays what pages the user segment has visited.

when currently viewing a report, click the *Update* button from the *Options* icon (⋮) in the top right corner.

Next, you'll discover how to use your user segments in a campaign.

64.3 Managing Campaigns

A campaign is an effort to expose certain user segments to a set of assets within a specific period of time. To manage campaigns for a Site,

1. Navigate to *Site Administration* → *Configuration* → *Audience Targeting* → *Campaigns*.
2. To create a new campaign, select the (+) button.
3. Set a *Name* and *Description.
4. Select the *User Segments* to target, a *Start Date*, an *End Date*, and a *Priority*.
5. Indicate if the campaign you create should be active or inactive.
6. Click *Save*.

The user segments you select when creating a campaign represent users targeted by the campaign. The start and end dates together specify the duration of the campaign. There can be multiple campaigns active at the same time that target the same user segments. In these situations, the priority attribute of the campaigns determines which campaign takes precedence. Finally, you can activate or deactivate a campaign via the *Active* attribute of a campaign. Deactivating a campaign disables the effect of the campaign. A deactivated campaign can be reactivated later. It can be useful to deactivate a campaign if a problem is found with the way content is being displayed. Once the problem has been corrected, the campaign can be reactivated.

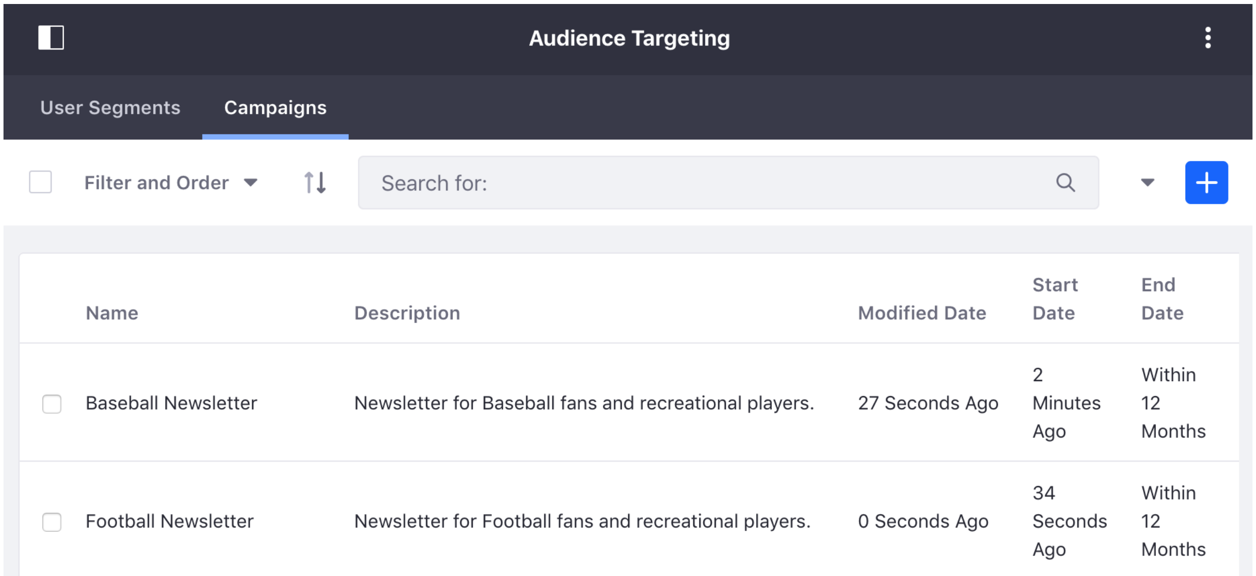

Figure 64.7: Navigate to Site Administration and click *Configuration* → *Audience Targeting* → *Campaigns* to manage campaigns for a Site.

For example, suppose you wanted to display certain content (life advertisements for your new Android app) to female baseball fans during the months leading up to the World Series. You could use the Gender rule (configured to female), the Device Rule (configured for Android devices), and the Score Points rule to define a user segment called *Female Baseball Fans*. The Score points rule assigns one point to a user each time the user visits a page or views an asset categorized under the user segment *Female Baseball Fans*. When a user accumulates a certain number of points (specified by the value of the Score Points rule's *Score Points Threshold* attribute), the user matches this rule. Once you have the user segment, you can create a new campaign targeting this segment, select start and end dates, choose a priority, choose *Active*, and then click *Save*. To present content to the users belonging to the *Female Baseball Fans* user segment, you must use the Campaign Content Display application.

Once you've created a campaign, you can open its summary view, which displays relevant data and configurations.

To configure permissions for creating, editing, and deleting user segments, visit the Roles section of the Control Panel. Then click on the *Actions* → *Define Permissions* button corresponding to the Role you want to configure permissions for, and search for *Audience Targeting*.

Managing Campaign Priorities

The priority of your campaigns becomes important when multiple campaigns are running at the same time on your website. The Campaign Content Display application can be configured to display

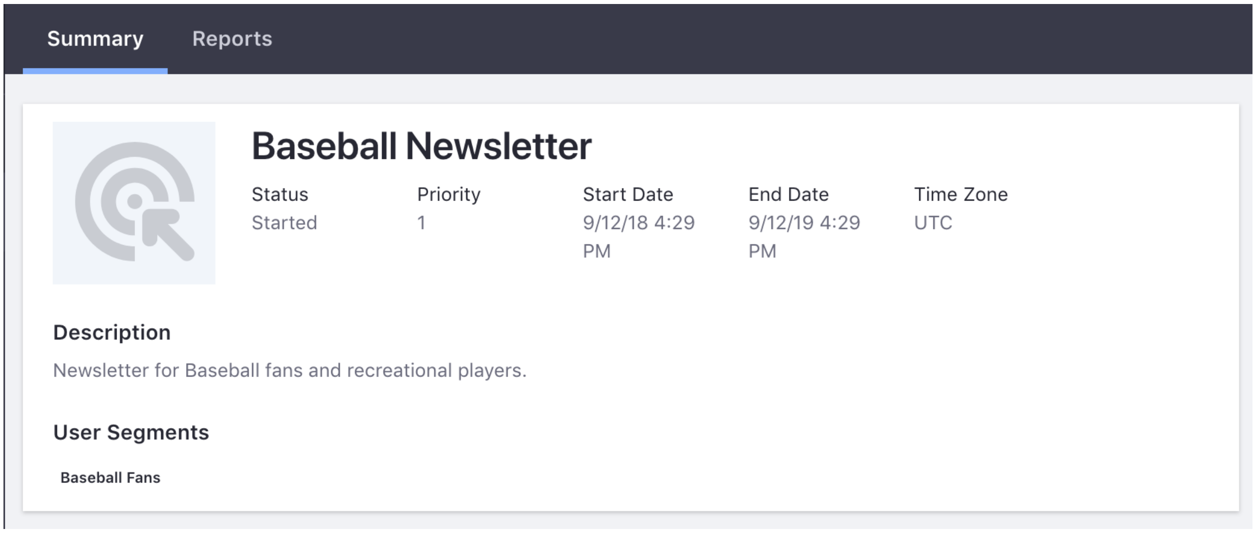

Figure 64.8: Select a pre-existing campaign to view its Summary page.

content based on the campaign your users match. When a user matches multiple campaigns, the one with the highest priority takes precedence.

If you have several Campaign Content Display applications around your website configured to display different content per campaign, changing the priority of one campaign automatically affects all Campaign Content Display applications. Similarly, if a campaign is deactivated or if a campaign's date range is exceeded, all Campaign Content Display applications on your website are affected.

64.4 Defining Metrics

To provide your marketing team with real feedback from users, you can define the user actions you want to track using *Metrics*. Metrics are used in reports to measure the effectiveness of a campaign by tracking certain user actions.

Metrics filter the analytics data gathered by the Audience Targeting Analytics engine to obtain the number of times a certain action was performed on a given element or content by users and their user segments.

Metrics are created by developers and deployed as extensions. Out of the box, Audience Targeting includes metrics to track the most common user actions. These metrics are described below.

Content

Tracks the number of times content has been viewed. Use the content selector to set the content to be tracked.

Page

Tracks the number of times a selected page has been viewed. You can track views on both both public or private pages.

Form

Tracks how many users view a form, interact with it (i.e., type or select values in the inputs), or submit it. If you select the *All* option from the *Event type* field, the custom report shows the figures for the three events simultaneously. You must also provide the form you want to track, which is selectable from the *Form* metric.

Link

Tracks how often links are clicked. This helps campaign administrators determine if they're sufficiently visible or helpful.

Similar to forms, you must provide the ID of the link you want to track. If you don't know it, you can inspect the HTML of the page where the link is and extract this information.

YouTube Videos

Tracks how users interact with embedded YouTube videos. You must enter the video's ID. You can extract this ID from the video URL as the value for the `v` parameter. For instance, in the URL `https://www.youtube.com/watch?v=2EPZxIC5ogU` the YouTube video ID is `2EPZxIC5ogU`. Then select one of the available events, or *All* to track all of them. For further reference on the meaning of these events, read the official YouTube API documentation.

Notice that this option only works if the YouTube video is embedded as an iframe. The iframe code is available from the YouTube video's *Share → Embed* menu.

Using Metrics

Suppose you want to run a campaign for an event that your company is hosting next month. You have created a main page for the event with a YouTube video and a *Register Now* banner. You also have a blog entry about the event displayed on several different pages and a Register page with the form to pay for registration. In this campaign, your goal is to get as many people to register as possible, but there is other information you want to track to ensure that everything is working as expected:

- Visits to the main page of the event
- Clicks to view the video
- Number of users who watched the video until the end
- Clicks on the Register Now banner
- Views of the blog entry about the event
- Views of the Register form
- Number of users who started to fill out the Register form
- Number of users who completed the registration

You can assign metrics to a campaign report, which is elaborated on in the next section. To access the Metrics palette,

1. Go to a pre-existing campaign.

2. Select the *Reports* tab.

3. Add a custom report.

   The Metrics palette is accessible at the bottom of the *New Report* wizard.

You could drag and drop *metrics* from the palette to track all the actions mentioned above. More types of metrics can be created by developers and deployed as OSGi plugins. See the Tracking User Actions with Audience Targeting tutorial for details.

Audience Targeting Analytics

Metrics uses the *Audience Targeting Analytics* engine that can be configured per site or per Liferay DXP installation.

1. To configure the analytics engine per Site, go to Site Administration and click *Configuration* → *Site Settings* → *Advanced* → *Audience Targeting Analytics*.

2. To configure it per portal instance, Go to *Control Panel* → *Configuration* → *System Settings* → *Audience Targeting Analytics*.

The following analytics options are available:

- Anonymous Users (not available per site)
- Pages
- Content
- Forms
  - Form Views
  - Form Interactions
  - Form Submits
- Links
- YouTube Videos

Tracking all the actions of all your users can be a heavy load for your server. Therefore, it's best to disable tracking any actions about which you don't need information. For example, by default, guest behavior analytics are tracked. This stores a large amount of data to the database. If you're not interested in tracking guest users,

1. Disable the *Anonymous Users* selector.

2. Click *Save*.

If you want to collect anonymous user data, but you are still mindful of database resources, you can change the interval for how often anonymous data is cleaned up. You can find the anonymous data storage setting under *Audience Targeting Service*.

- **Check Interval** and **Check Time Unit** define how often data cleanup takes place. If interval is set to *1* and unit is set to *Day* then the clean up task occurs once per day.

- **Max Age** and **Max Age Time Unit** define how old the data must be to be deleted, so that data created immediately before the cleanup task is not immediately deleted. If the age was set to *10* and the unit set to *Hour*, any data older than 10 hours is removed, and any data less than 10 hours old is preserved until the next cleanup.

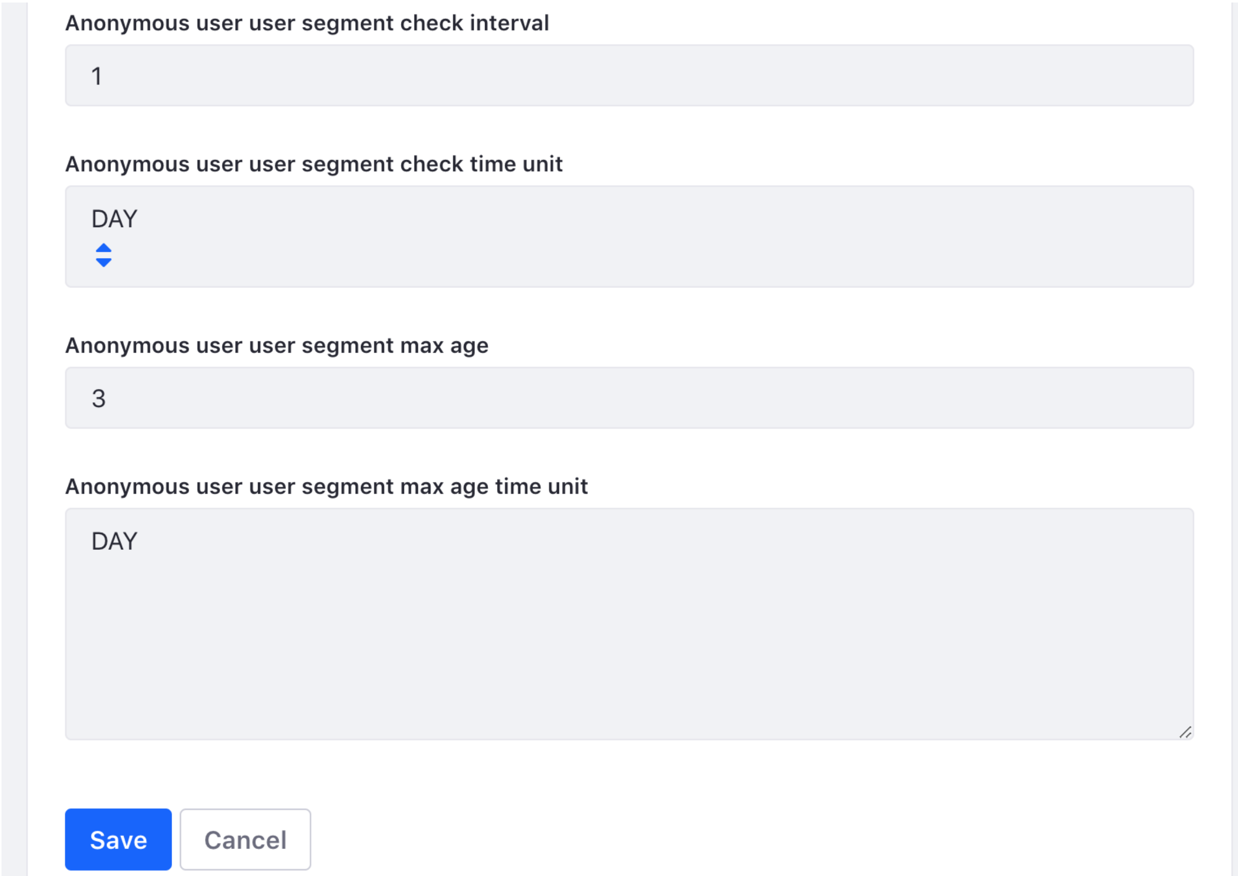

Figure 64.9: You can manage anonymous data cleanup here.

Disabling analytics for certain entities means they aren't tracked. Carefully manage analytics to optimize your Audience Targeting experience.

You can also store your analytics data in a separate database schema, which allows for independent scalability. To separate the storage of analytics data from Liferay's database schema,

1. Navigate to the Control Panel → *Configuration* → *System Settings* → *Web Experience*
2. Select *Audience Targeting Analytics Storage*.
3. Fill out the external storage fields to point to your alternative database schema.

Once you've saved your external datasource configuration, you must restart the Audience Targeting Analytics component.

1. Navigate to the Control Panel → *Apps* → *App Manager* and select the *Liferay Audience Targeting* app suite.

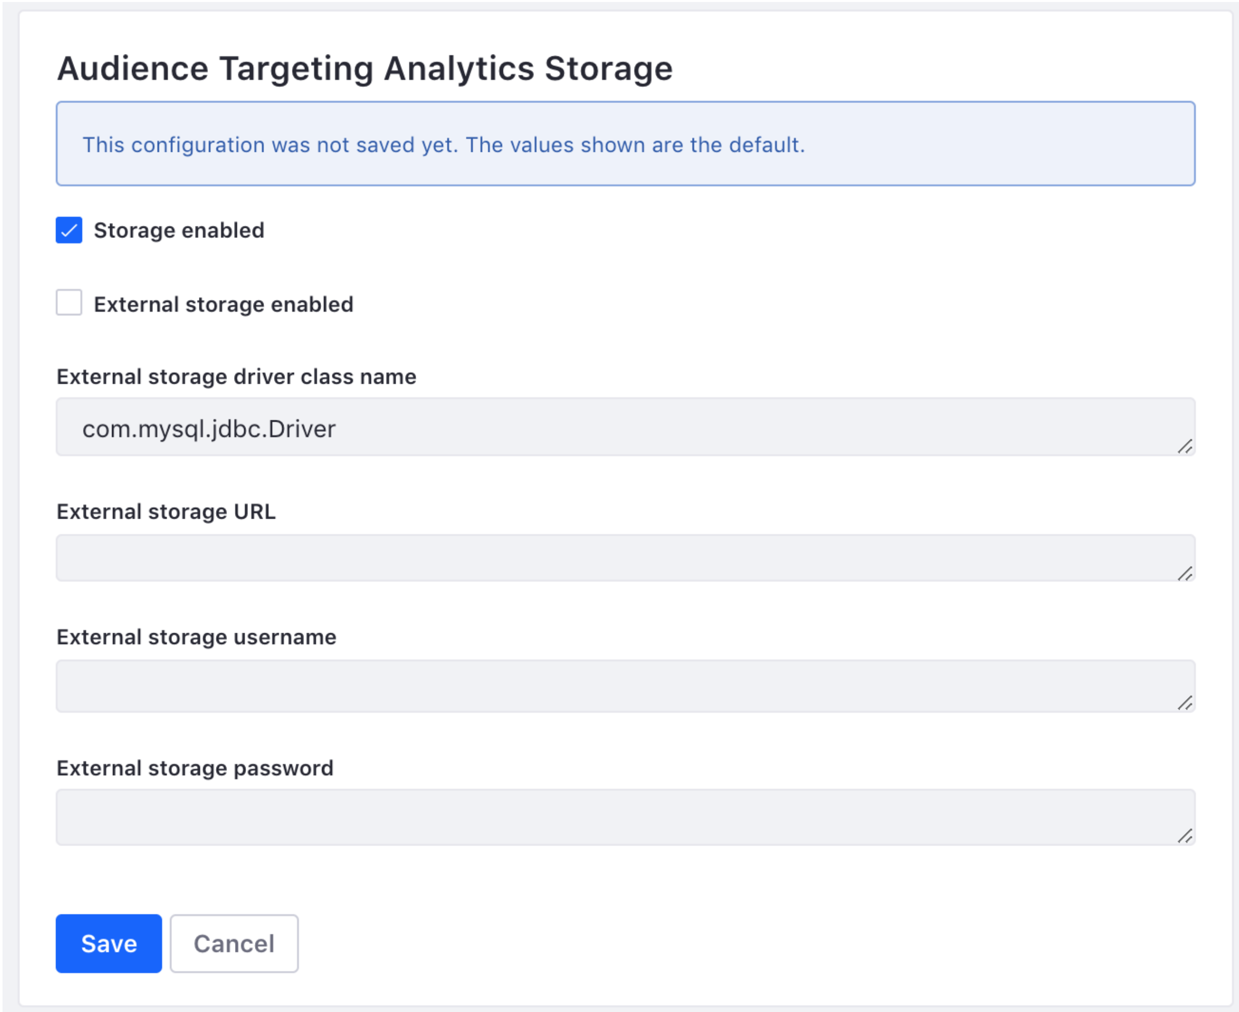

Figure 64.10: By filling out the external storage requirements, you configure your Audience Targeting analytics data to be stored in an alternative database schema.

2. Select the *Options* ( ) button for the Analytics component and click *Deactivate*.

3. Select the *Options* ( ) button for the Analytics component again and click *Activate*.

64.5 Campaign Reports

Campaigns are useless without reports showing what happened. Select the campaign name and click the *Reports* tab to see the list of reports available. To create a report,

1. Click on a *Campaign* name.

2. Open the *Reports* tab.

3. Click *Add Custom Report* (+).

The Content Views report shows the number of times assets were viewed via the Campaign Content Display application by users in the context of the current campaign. For example, if you placed three Campaign Content Display widgets on your website to display content for a campaign, the Content View report would show how many times users interacted with that content.

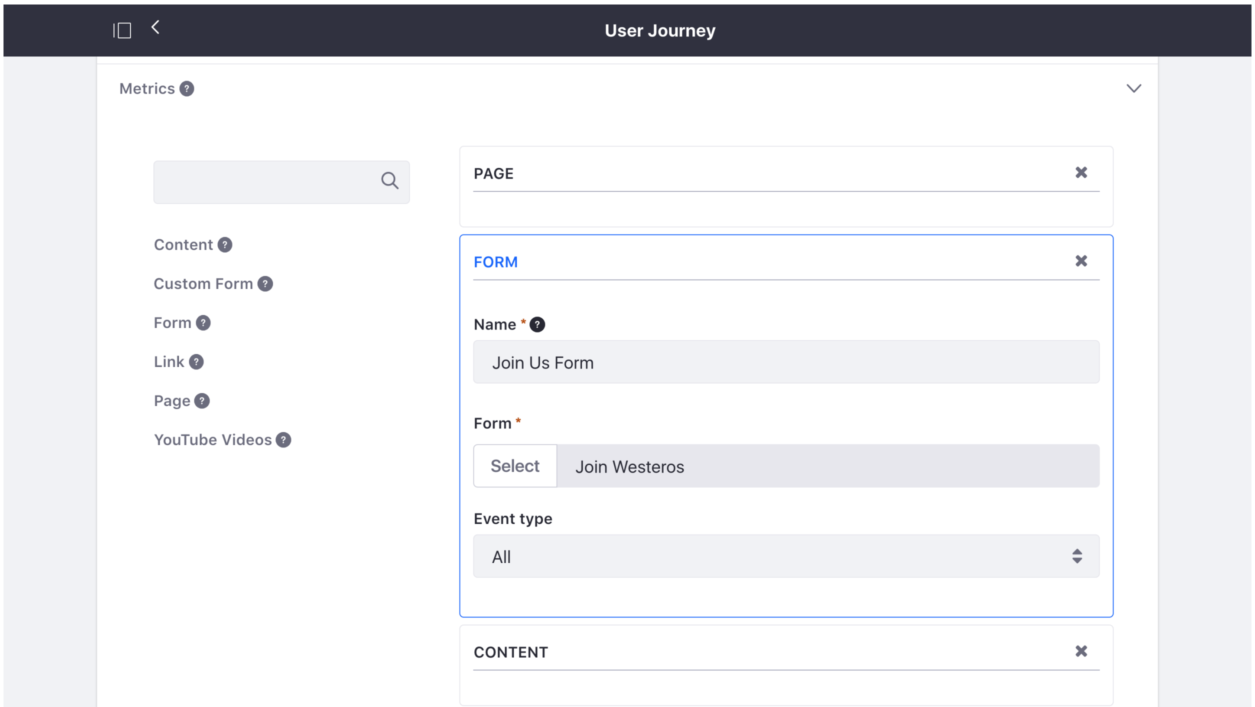

Figure 64.11: You can build your own custom campaign report to fit your needs.

You can track many other user actions by creating a custom report. You can drag and drop different kinds of metrics to track for your campaign, and the report shows how many times each metric was triggered by users.

Consider the metrics example we introduced earlier: you've created a campaign for a future event. For this event, you created a main page with a YouTube video and a banner stating "Register Now." You also have a blog post about the event that appears on several different pages. Finally, the Register page contains the form to pay for the event. A custom Metrics report could show you how many users visited the event page, how many watched the video, how many clicked on the banner, how many viewed the blog post about the event, how many started filling the registration form, etc. This information helps you measure the effectiveness of your campaign. You can use this information to evaluate if users are following the engagement path you had prepared.

The metrics you apply to a report can be tracked in several different ways. For instance, if you added the *YouTube Videos* metric to your report, there are several different event types to track, such as how many users played the video or who finished the video. You can track these events with the *Event type* field. You can even track *all* events, if you want all the metrics for an option.

64.6 Using the Audience Targeting Widgets

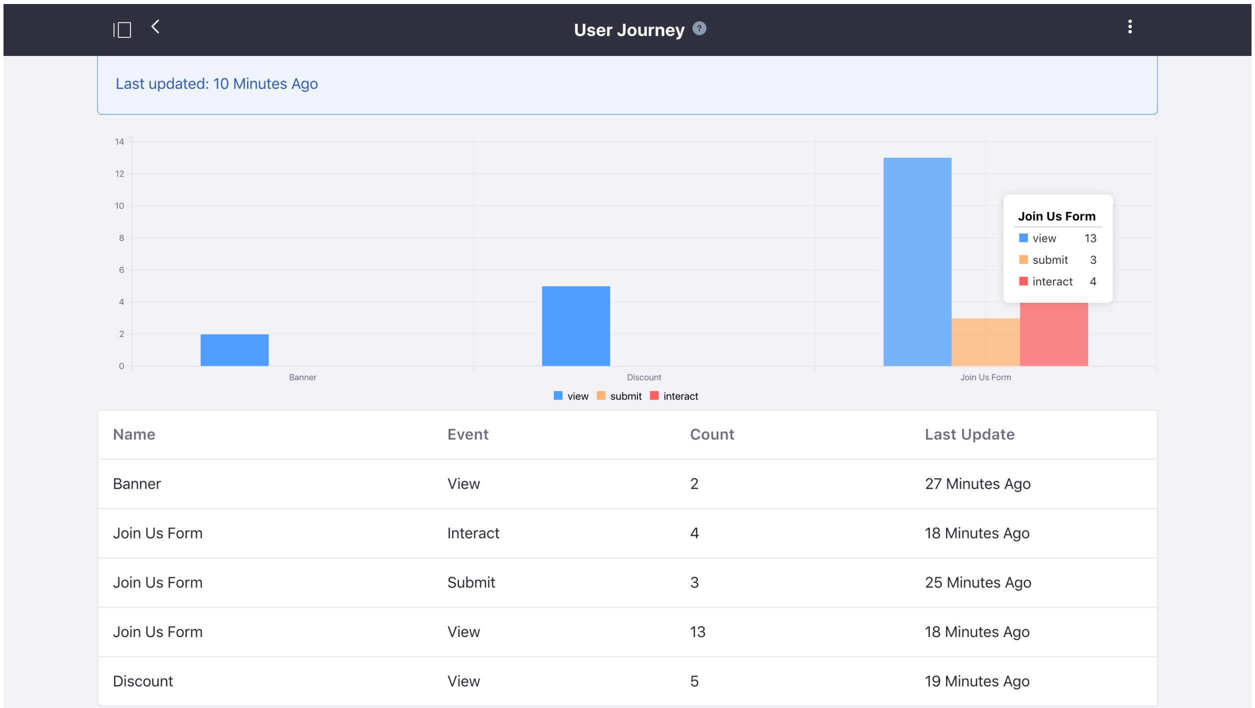

Figure 64.12: This campaign report displays several event types for content in the campaign.

Audience Targeting includes widgets that can be added to any portal page:

- Campaign Content Display
- User Segment Content Display
- User Segment Content List

All of these widgets support Application Display Templates (ADTs), and any of Liferay DXP's out-of-the-box Asset Publisher ADTs can be re-used for these Audience Targeting widgets.

User Segment Content Display

User Segment Content Display displays content to users based on their user segment membership. Administrators can specify multiple rules according to the following format:

- If the user `[belongs|does not belong]` to `[any|all]` of the following user segments `[specify a list of user segments]`, then display this content: `[specify a specific asset]`.

You can specify any number of *if* clauses when configuring the User Segment Content Display widget's rules. However, an *otherwise* clause always follows the last *if* clause so the widget knows what to display if the user doesn't match any user segments. *Don't display anything* is an option for the *otherwise* clause.

For example, you can add a User Segment Content Display widget to a page and configure the following rules for it:

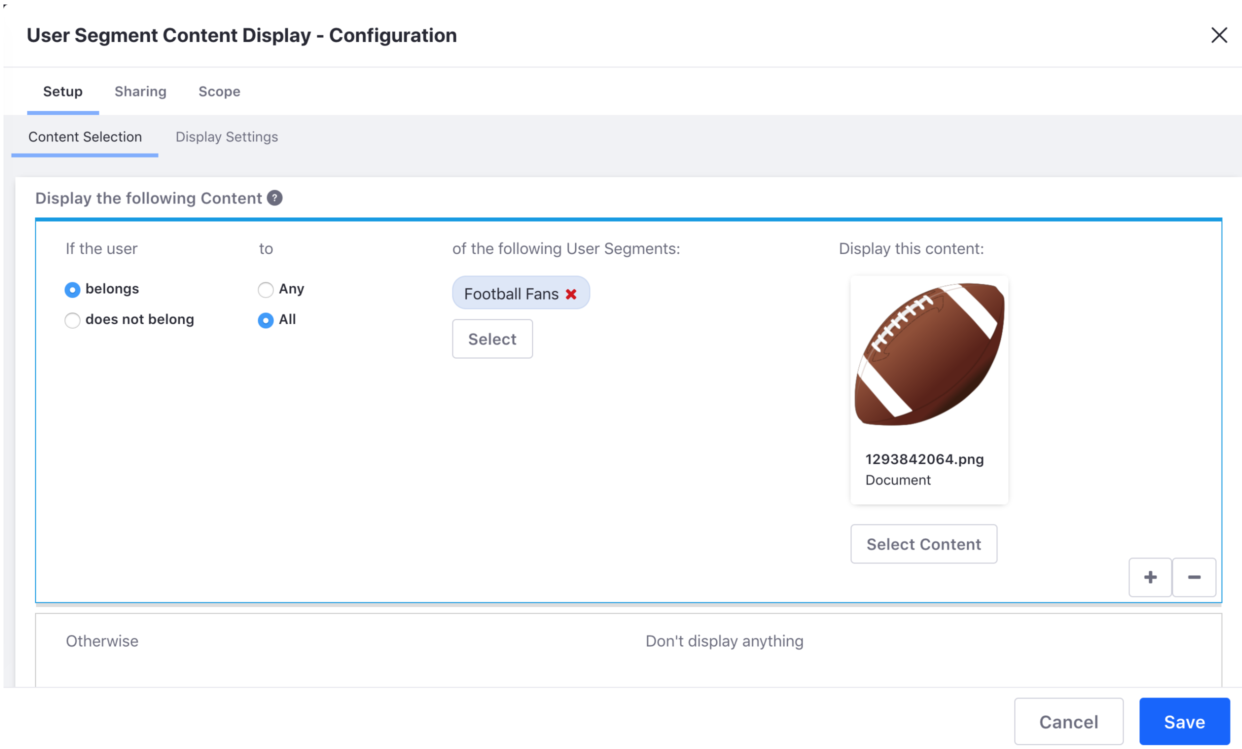

Figure 64.13: You can configure the User Segment Content Display widget to display content according to rules that you define in the widget's configuration window.

- If the user *belongs* to *any* of the following user segments: *Tennis fans*, display this content: *tennis_picture.jpg*.
- If the user *belongs* to *any* of the following user segments: *Basketball fans*, display this content: *basketball_picture.jpg*.
- If the user *belongs* to *any* of the following user segments: *Soccer*, display this content: *soccer_picture.jpg*.
- Otherwise, *Don't display anything*.

Once the widget has been configured this way, users (even guest users) see a different image based on their user segment. The widget isn't visible to users who don't belong to any of the configured user segments.

Administrators can preview the various assets that have been configured for different user segments.

Campaign Content Display

Campaign Content Display is similar to the User Segment Content Display except that instead of displaying assets based on user segment membership, it displays an asset based on the campaigns that a user matches. However, the Campaign Content Display widget's display rules are simpler than those of the User Segment Content Display widget. You can specify multiple rules for the Campaign Content Display widget according to the following format:

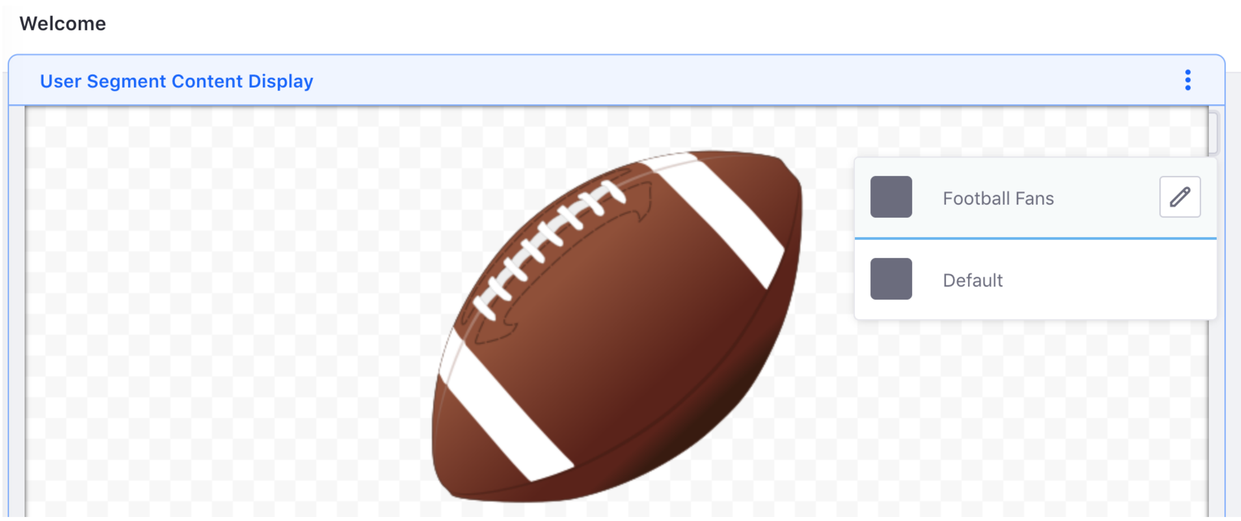

Figure 64.14: In the User Segment Content Display widget, site administrators can preview the various assets that have been configured to be displayed to different user segments.

- If the user belongs to this campaign: `[select a campaign]`, then display this content: `[specify a specific asset]`.

As with the User Segment Content Display, you can specify any number of *if* clauses when configuring the Campaign Content Display widget. An *otherwise* clause always follows the last *if*. These rules cannot be ordered manually as they can with the User Segment Content Display widget. The order of the rules is based on the priority of the campaigns.

For example, if you created two campaigns called *World Cup* and *Stanley Cup* that targets users interested in the World Cup of soccer and the Stanley Cup of hockey, you could add a Campaign Content Display widget to a page and configure it with the following rules:

- If the user belongs to this campaign: *World Cup*, display this content: *soccer-ball.jpg*.
- If the user belongs to this campaign: *Stanley Cup*, display this content: *hockey-.png*.
- Otherwise, display this content: *generic-sports-jersey.jpg*

Once a Campaign Content Display has been added to a page and been configured this way, users matching the World Cup campaign see the soccer image. Users matching the Stanley cup campaign see the hockey image. Users who don't match either see the default image. Of course, once a campaign has ended, no users match that campaign. Once all campaigns have ended, all users see the default image.

Note: When Audience Targeting is enabled with a large number of rules, tracking actions, and reports, you may need to adjust some JVM parameters of your server for optimal performance.

The Campaign Content Display widget, like the User Segment Content Display widget, allows Site administrators to preview the different assets that are displayed for different campaigns.

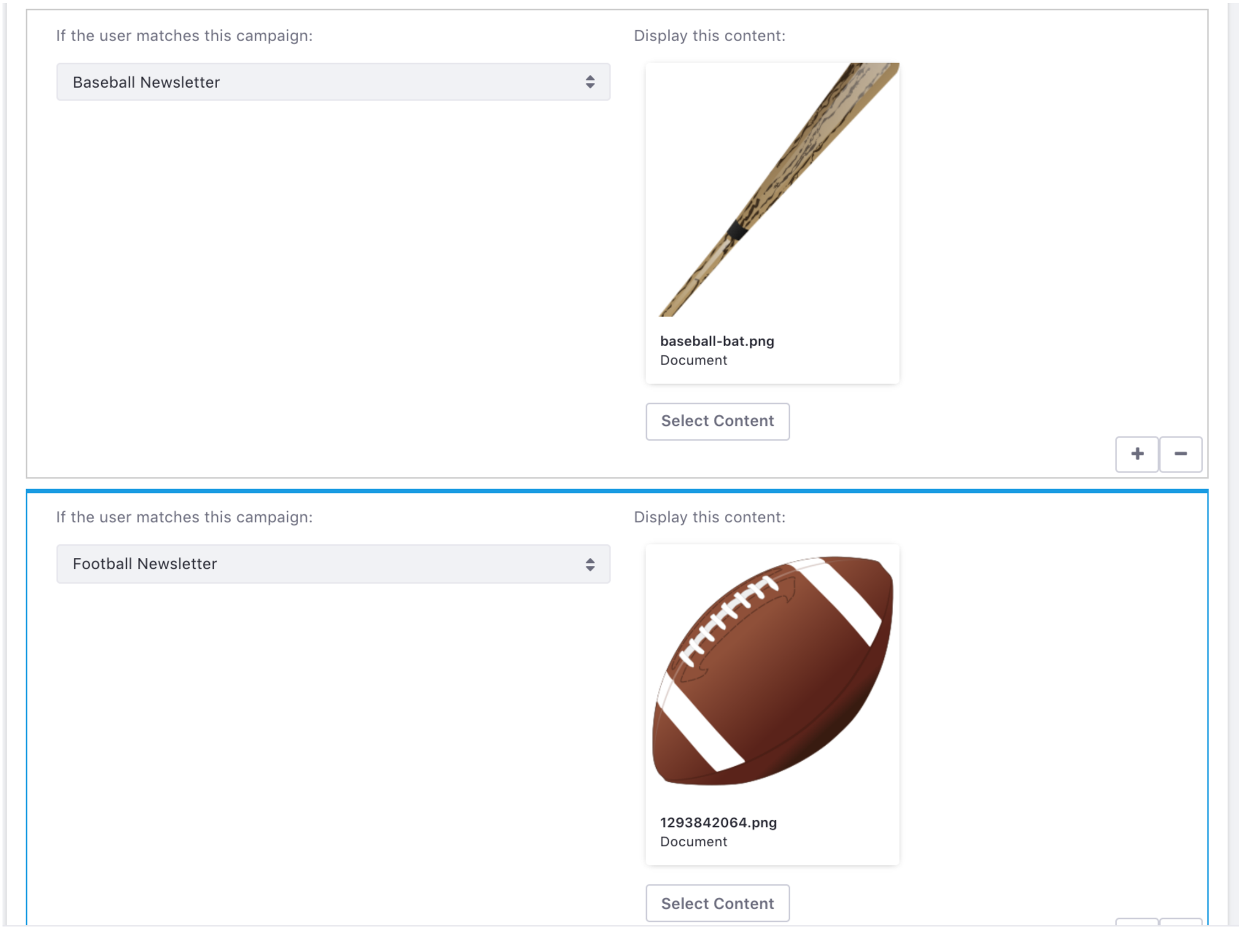

Figure 64.15: The rules for configuring the Campaign Content Display widget to display content are similar to the rules of the User Segment Content Display widget, but simpler.

Asset Publisher

For 7.0 the Asset Publisher replaces the functionality of the User Segment Content List. When Audience Targeting is installed, the Asset Publisher gains an additional configuration option: *User Segments Filter*. When *User Segments Filter* is activated, the only assets displayed are those that match the filters configured in the widget and are categorized under any of the user segments for the current user. Assets not categorized with User Segments don't appear in any Asset Publisher with the *User Segments Filter* enabled. The Asset Publisher can be configured to display these assets in a list or to display the full content.

Next you'll learn how to simulate your user segments and campaigns.

64.7 Simulating User Segments and Campaigns

Once you create all those segments, you'll likely want to test them. The Audience Targeting Simulator can preview the way pages appear to different user segments. It appears under the *Simulation* icon ( ) in the top Control Menu.

1. Click on the Simulator icon to open the right-side panel.
2. Select *User Segments*.

The Audience Targeting Simulator is only visible to users with the View permission for the Audience Targeting Simulator.

Note: You can grant Audience Targeting permissions by clicking *Actions* → *Define Permissions* menu for the appropriate Roles in the Control Panel. When defining permissions for a Role, go to *Site Administration* → *Configuration* to find the *Audience Targeting* section.

Go to the User Segments tab of the Audience Targeting simulator to see a list of all the user segments available for the Site. Selecting segments lets you see how the website would appear to them. You can navigate around your Site as if you were a user that matched the selected user segments.

Next, you'll examine all the Audience Targeting rules available for you to use by default.

64.8 Audience Targeting System Settings

Audience Targeting's configuration options are provided in System Settings. Configurations made here take effect system-wide. You can find Audience Targeting system settings by navigating to the Control Panel → *Configuration* → *System Settings* → *Audience Targeting*. The following options are available for the Audience Targeting apps:

- Audience Targeting Analytics Service
- Audience Targeting Analytics Storage
- Audience Targeting Service
- Campaign Content Report
- Campaign Tracking Actions Report
- User Segment Content Report

The settings described here are unique to the Audience Targeting application. Some system setting groups may have more configurations than what are defined here; unspecified settings are generic configurations that are applied across many Liferay DXP applications.

Audience Targeting Analytics Service

These settings define the general behavior of the Audience Targeting Analytics modules.

Analytics events check interval: Defines the interval (in days) between data check events.

Analytics events max age: Defines the max age (in days) of events stored in the analytics tables. When a check event is triggered, entries older than the max age are removed.

Analytics flush interval: Defines the interval (in ms) between flushes of the analytics event buffer on the client-side.

Get analytics key: Currently only used for testing purposes.

Content tracking enabled: Enables Audience Targeting.

Form excluded ids reg exp: Defines forms' excluded IDs using regular expressions.

Form tracking enabled: Enables tracking capabilities for forms.

Form interaction tracking enabled: Enables tracking capabilities for interactions.

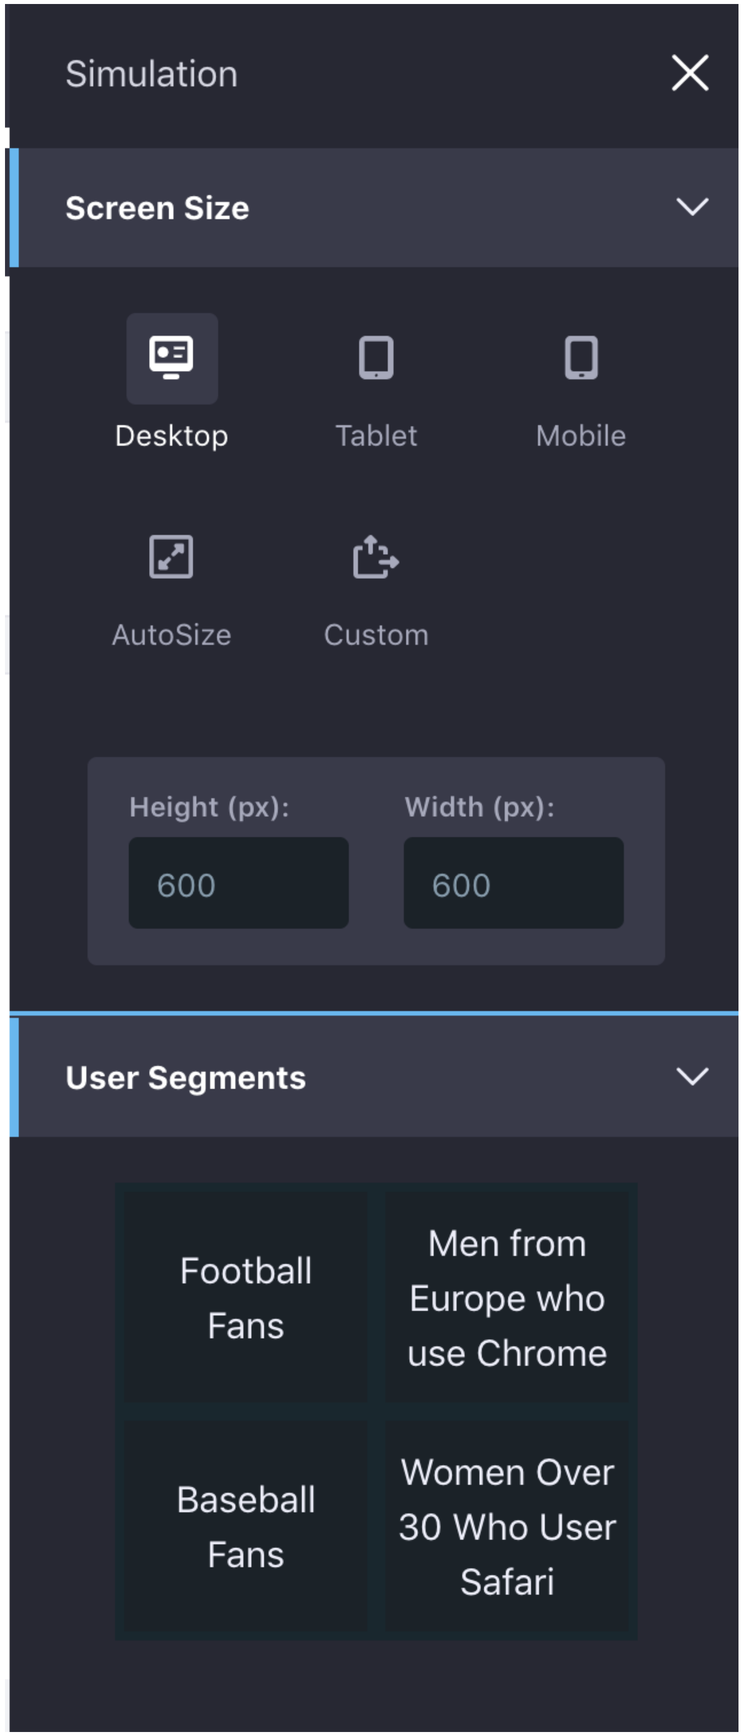

Figure 64.16: Audience Targeting provides a simulator in the Simulation menu.

Form submit tracking enabled: Enables tracking capabilities for form submissions.
Form view tracking enabled: Enables tracking capabilities for form views.
Link excluded ids reg exp: Defines links' excluded IDs using regular expressions.
Link tracking enabled: Enables tracking capabilities for links.
Link click tracking enabled: Enables tracking capabilities for clicking events on links.
Page tracking enabled: Enables tracking capabilities for pages.
YouTube tracking enabled: Enables tracking capabilities for YouTube videos.

Audience Targeting Analytics Storage

These settings configure the analytics engine used by Audience Targeting.
External storage enabled: Enables analytics storage.
External storage driver class name: Defines the external driver class name for analytics storage.
External storage URL: Defines the external storage URL for analytics storage.
External storage username: Defines the external storage user name for analytics storage.
External storage password: Defines the external storage password for analytics storage.

Audience Targeting Service

These settings define the cleaning task for the record of user assignments (anonymous and registered) to user segments (i.e. the `AnonymousUserUserSegment` table).

Anonymous user user segment check interval: Defines the interval (in days) between data check events.

Anonymous user user segment max age: Defines the max age (in days) of records stored in the `AnonymousUserUserSegment` table. When a check event is triggered, entries older than the max age are removed.

Campaign Content Report

These settings configure campaign content reports.

Campaign content report check interval: Defines the report check interval (in hours) for campaign content. Since the lifetime of events in the Analytics tables is limited to the *Analytics events max age*, reports must retrieve and process the latest events from Analytics with a shorter time interval.

Campaign Tracking Actions Report

These settings configure campaign tracking action reports.

Campaign tracking action report check interval: Defines the report check interval (in hours) for campaign tracking actions. Since the lifetime of events in the Analytics tables is limited to the *Analytics events max age*, reports must retrieve and process the latest events from Analytics with a shorter time interval.

User Segment Content Report

These settings configure user segment content reports.

User segment content report check interval: Defines the report check interval (in hours) for user segment content. Since the lifetime of events in the Analytics tables is limited to the *Analytics events max age*, reports must retrieve and process the latest events from Analytics with a shorter time interval.

64.9 Liferay Audience Targeting Rules

Rules enable Audience Targeting administrators to define how users are classified into the user segments they create.

A rule evaluates if the current user matches or not one or several conditions. Rules can be configured and combined differently for each user segment to create a specific classification algorithm. A user must match all rules in a user segment to be classified into it.

Behavior Rules

These rules classify users based on their navigation behavior.

These rules don't work properly if content tracking or page tracking is not enabled. For Site level configuration,

1. Go to *Site Administration* → *Configuration*
2. Select *Site Settings* and then the *Advanced* tab.
3. Find the *Audience Targeting Analytics* heading.

For portal level configuration,

1. Go to *Control Panel* → *Configuration* →
2. Select *Instance Settings* and the *Audience Targeting Analytics*.

Content Visited

Evaluates if the user has visited the selected content.

Page Visited

Evaluates if the user has visited the selected page.

Previously Visited Site

Evaluates if the user has visited a specific Site before visiting the current Site. The Site's URL can be evaluated based on whether it contains or starts with defined text.

For example, if you want to segment users based on whether they navigated to your Site from Twitter, you could define a Twitter URL for this rule.

Score Points

Evaluates if the user has visited any page or content categorized under this user segment a number of times equals or higher than the configured threshold.

Social Rules

Social rules classify users based on their profiles in social networks (e.g., Facebook).

These rules do not work properly if login through the specific social network (Single Sign On) is not enabled and properly configured. You can learn more about using Single Sign On in the Facebook Connect SSO.

Facebook Age

Evaluates if the age in the user's Facebook profile is in the configured range.

Facebook City

Evaluates if the city in the user's Facebook profile matches the configured value.

Facebook Education

Evaluates if the education in the user's Facebook profile matches the configured values.

Facebook Friends

Evaluates if the user has more or less friends in Facebook than the selected value.

Facebook Gender

Evaluates if the gender in the user's Facebook profile matches the selected value.

Facebook Page Like

Evaluates if the user likes the configured Facebook page.

User Attributes Rules

These rules classify users based on the attributes.

These rules do not work properly if the corresponding attributes from the user profile are not available. Follow the given instructions for each rule in such cases.

Age

Evaluates if the age in the user's profile is in the configured range.

Custom Field

Evaluates if the value set for the selected custom field matches the value from the user's profile.

Gender

Evaluates if the gender in the user's profile matches the selected value.

Organization Member

Evaluates if the user is member of the selected Organization.

Organization Role

Evaluates if the user has the selected Role assigned in the selected organization.

Regular Role

Evaluates if the user has the selected regular Role assigned.

Site Member

Evaluates if the user is a member of the selected Site.

Site Role

Evaluates if the user has the selected Role assigned in the selected Site.

User Group Member

Evaluates if the user is member of the selected User Group.

User Signed In

Evaluates if the user is signed in.

Session Attributes Rules

These rules classify users based on their session attributes. Session attributes come from Java `HttpSession` data.

These rules do not work properly if the corresponding session attributes are retrieved through any tools that are not currently available or properly configured. Follow the given instructions for such rules in such cases.

Access Time

Evaluates if the user visited at a time that is within the configured range. Notice that the reference time is that of the server.

Browser

Evaluates if the user visited with the selected browser.

Device

Evaluates if the user visited with the selected device.

This rule is based on the existing Device Families. To manage Device Families, go to *Site Administration → Configuration → Mobile Device Families*.

IP Range

Evaluates if the user visited with an IP with the configured range.

Language

Evaluates if the user visited with the selected language.

Last Login Date

Evaluates if the last login date of the current user was after, before, or between the selected dates.

Location

Evaluates if the user visited from a country and region that match the configured values.

Operating System

Evaluates if the user visited with the selected Operating System.

Sign Up Date

Evaluates if the Sign Up date of the current user was after, before, or between the selected dates.

Liferay's Audience Targeting app can define custom user segments, target specific content to different user segments, and create campaigns for different user segments. You can generate reports to measure the effectiveness of your campaigns. A simulator previews how your Site appears to users belonging to different user segments. You can create user segments by applying various rules that reference session attributes, profile attributes, behavior, and other information, such as information from social networks. If you need more, developers can create additional rules and rule types.

Chapter 65

Managing Users

Ever heard a retailer advertise as a "one stop shop" for anything you want? The idea is they have so much stuff that whatever you want is probably there. Liferay's Control Panel is like this. Where do you create users, organizations, or sites? Where do you configure permissions and plugins and pretty much anything else? You do it from the Control Panel. When signed in as an administrator, you can access the Control Panel from the Product Menu by clicking *Control Panel*.

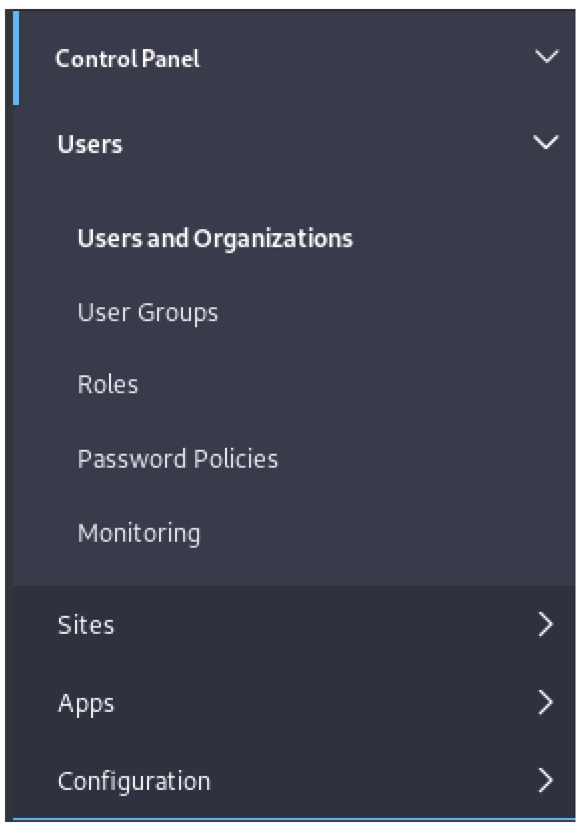

Figure 65.1: Administrators can access the Control Panel from the Product Menu.

The Control Panel is divided into four main areas: Users, Sites, Apps, and Configuration. The

Users section lets you create and manage users, organizations, user groups, roles, and password policies. If monitoring has been enabled for your portal, you can also view all the live portal sessions of your users.

Anonymous User: *Anonymous Anonymous* is used for the new Managing User Data functionality. Created the first time an administrator clicks *Delete Personal Data* for a User, *Anonymous Anonymous* is a deactivated User assigned anonymized assets. The Anonymous User is configurable, so the name and configuration details might be different in your virtual instance.

65.1 What's New with User Management?

In Liferay DXP version 7.0, the user management functionality and behavior didn't change much from previous versions. Improvements to existing behavior were introduced, but most of the work went into the back-end code so that future enhancements could be made. In 7.0, we're excited to announce some new features for user administrators:

1. Numeric screen names can be used without issue out of the box. Read here for the details.
2. Add organization types from System Settings. There's no need for portal properties anymore, or the requisite system restart. Read here for the details.
3. Configure default user profile pictures from System Settings. Read here for the details.

Begin exploring Liferay's User Management functionality by reading about adding and editing users.

Chapter 66

Users and Organizations

Users and *Organizations* are fundamental entities. If your portal requires people (even just a set of site administrators) to have accounts to do anything, you need to know about users. If your users are at all divided hierarchically, like into departments, you'll find that organizations are helpful.

You're probably not surprised to hear that Users and Organizations are managed in the Control Panel's *Users and Organizations* section. If it were any different, it'd be weird.

Consider the Lunar Resort portal. Consider what you'd do if

- An employee leaves the company to join that pesky competitor, Martian Resort and Luxury Spa.
- An employee joins the resort as a new Mechanical Crew member.
- An employee is promoted from Crew Supervisor to Department Head and needs the requisite permissions.
- You need to organize the users by department.
- A new department is added to the Lunar Resort and the employees need their own internal website.
- An employee gets married, and their name changes.

The user tasks listed above are all resolved in the Users and Organizations section of the Control Panel.

66.1 What are Users?

In case there's any confusion over the term, a User is an entity that can sign into the portal and do something. Generally a User has more privileges, called Permissions, than a Guest of your site, who does not sign in. Users are assigned Roles, and Roles define the User's privileges.

Understanding Users is pretty straightforward. Organizations are a bit trickier, but a smart administrator like you is undoubtedly up to the challenge. Read more about Organizations here.

The remaining articles in this section give you guidance on managing (creating, deleting, editing, and more) Users and Organizations.

66.2 Adding, Editing, and Deleting Users

At the root of managing Users is adding, editing, and deleting them. As long as you're the Administrative user, you can do all these things and more.

Adding Users

Adding Users is straightforward:

1. From the Product Menu, click *Control Panel* → *Users* → *Users and Organizations*.
2. In the Users tab, click the *Add* button .

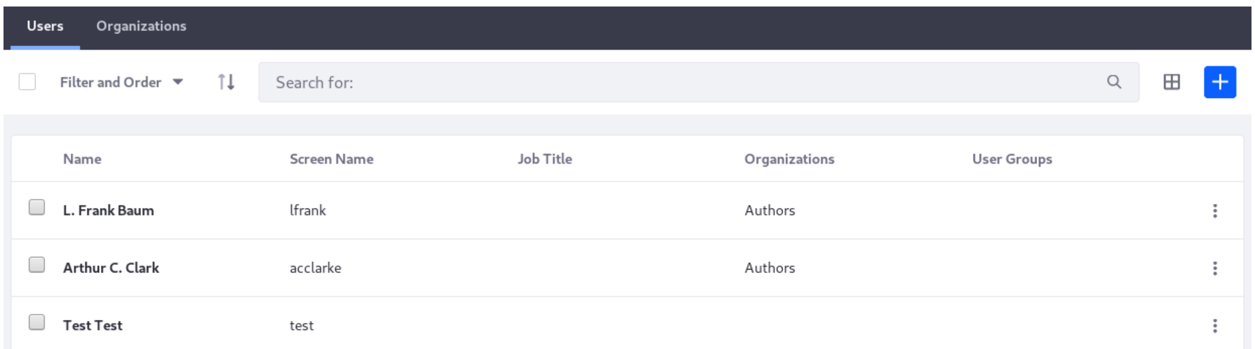

Figure 66.1: Add Users from the Users and Organizations section of the Control Panel.

3. Fill out the Add User form and click *Save*. At a minimum, provide a Screen Name, First Name, Last Name, and Email Address for the User.

```
Note: Screen names and email addresses are not interchangeable. A screen
name cannot contain an `@` symbol because it is used in the URL to a user's
private page.
```

```
The Add User functionality is split over several independent forms. Saving
the first form creates the user, and then you'll see a success message
saying

    Success. Your request completed successfully.
```

After submission of the first form, you see a larger form with many sections. The one you're on is the Information section. To the left is a navigation pane where you can continue configuring the user you're adding by clicking through the available sections. The options in the left menu change as you click through the tabs at the top. Peruse the sections for the three tabs (General, Contact, Preferences) and fill in all the applicable information.

You don't have to fill anything else out right now. Just note that when the user account was created, a password was automatically generated. If Liferay was correctly installed and a mail

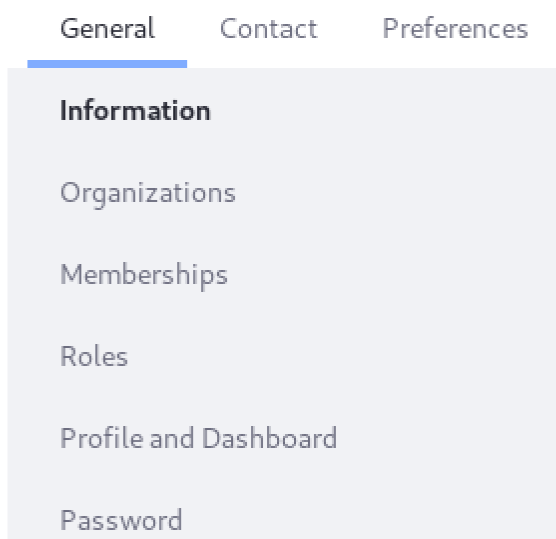

Figure 66.2: At a minimum, enter a screen name, email address, and first name to create a new user account. Then you'll be taken to the Information form and can continue configuring the user.

server was set up, an email message with the User's new password was sent to the User's email address.

If you haven't set up a mail server, click the *Password* item from the General menu and manually set a password for your new user. Enter the new password twice.

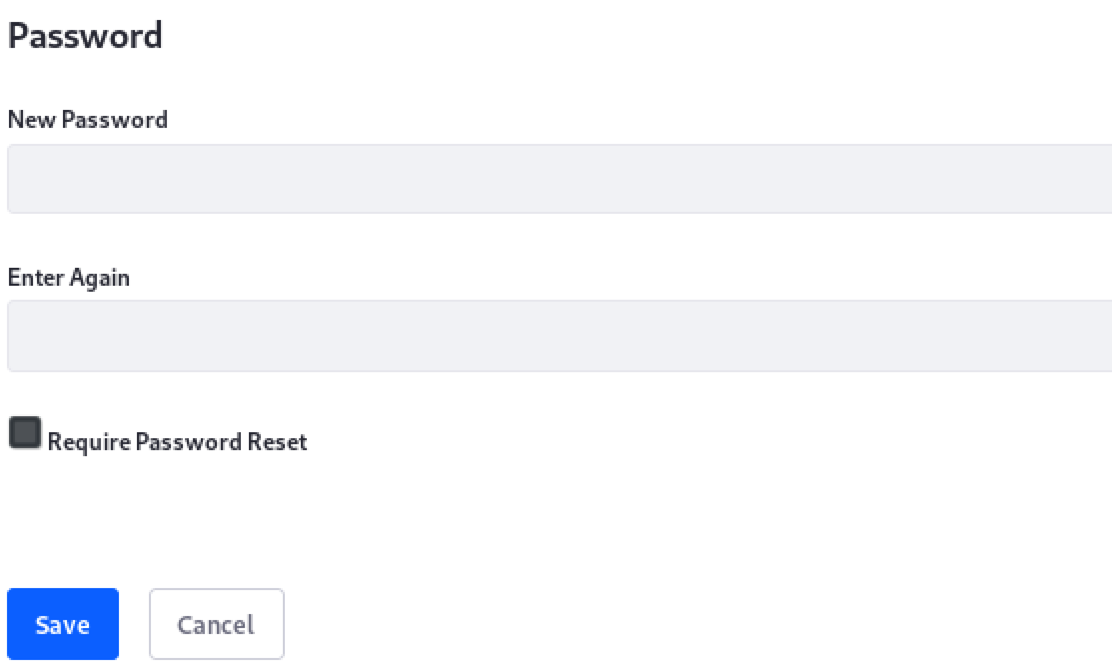

Figure 66.3: Enter the password twice to manually set the password for a user. If the Password Policy you're using is configured to allow it, select whether to require the user to reset their password the first time they sign in to the portal.

Editing Users

If you click on *Users and Organizations* in the Control Panel, you'll see your own user's account in the list of Users, along with any others. To change something about a particular user, click the *Actions* button (⋮) next to that user.

Choosing *Edit* takes you back to the Edit User page where you can modify any aspect of the User account including the screen name, email address, first name, last name, Site and Organization memberships, Roles, etc.

Choosing *Permissions* allows you to define which Roles have permissions to edit the User.

Choosing *Manage Pages* allows you to configure the personal pages of a User.

Choosing *Impersonate User* opens another browser window that loads the site as if you were the User so you can test your User management on a User to make sure you're achieving the desired behavior, without having to repeatedly log out of your administrator account and into the User's account.

Choosing *Deactivate* deactivates the user's account. The User is still in your database along with all the rest of your Users, but the account is deactivated, so the User cannot sign in to the portal. You can toggle between active and inactive Users in the Users view. If all the Users are active, this filtering option doesn't appear.

Choosing *Erase Personal Data* deletes the User's personal data.

Choosing *Export Personal Data* lets you download the User's personal data.

Most Users can't perform any of the above actions. In fact, most Users won't have access to the Control Panel at all. You can perform all of the above functions because you have administrative access.

Deleting Users

You must be careful when deleting Users. To guard against accidental deletion of Users, a two-step process must be followed: deactivate first, then delete.

1. Find the User to delete in the Users tab of *Control Panel → Users → Users and Organizations*. If you have a lot of Users, save time by searching for the User.

2. Click the *Actions* menu for the User and select *Deactivate*. You're asked to confirm that you want to deactivate the User. Click *OK*.

   You'll see a success message and the User disappears, but isn't gone yet.

3. By default the Users table displays only Active users. Click on *Filter and order* in the top of the table and a dropdown menu appears. Click *Inactive*, and you can see the User you just deactivated.

4. Click the Actions menu again, and click *Delete* if you really mean to delete the User. Confirm that you want to delete the User, and now the User is gone. This time, it's for real.

Deactivated Users: Deactivating a user means the user can't log in to the portal. He/she has no more permissions in the sites and pages of the portal than a guest, although the account still exists in the system.

Users are reactivated by a portal administrator finding them in the Users table (be sure you're filtering the table results by Deactivated users), clicking the Actions menu, and selecting Activate.

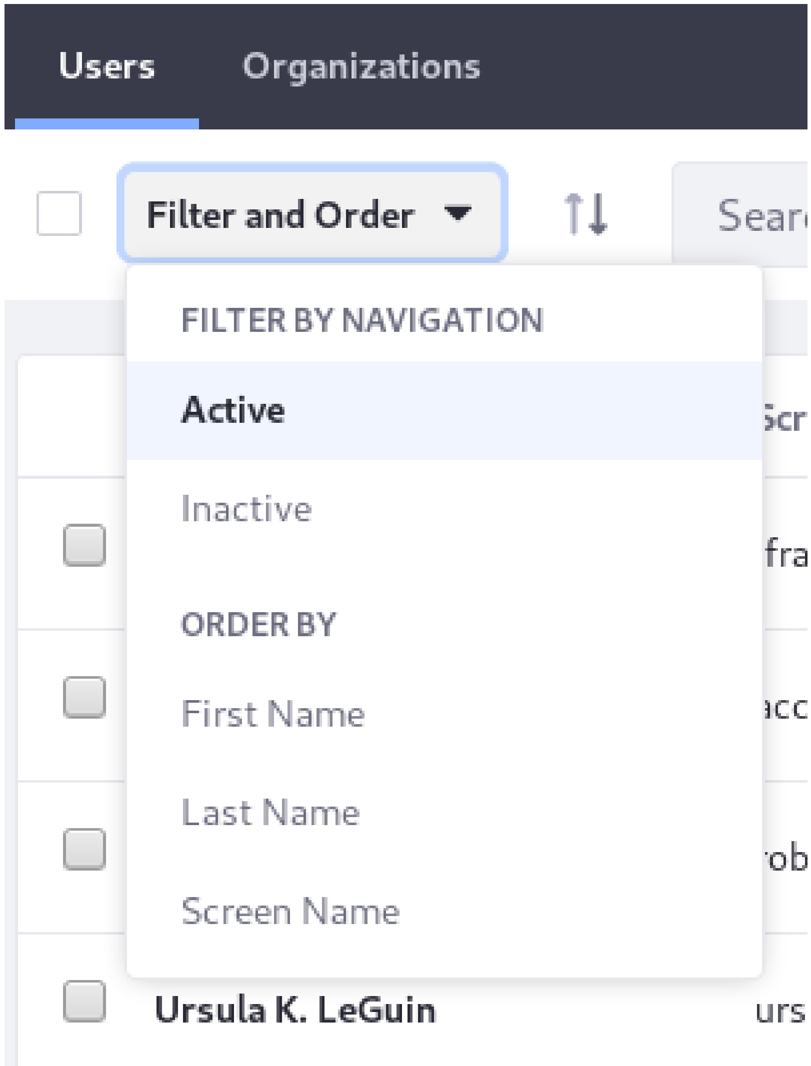

Figure 66.4: You can choose whether to view active or inactive (deactivated) portal users in the users list found at *Product Menu → Control Panel → Users → Users and Organizations*.

There's no confirmation window for activation: they're automatically restored to their former status once Activate is clicked.

Now you understand the basic principles of user administration. There are important additional topics in the next article that you should consider mandatory information for all portal administrators, so do continue reading.

66.3 User Management: Additional Topics

You've learned the basics on adding and editing Users, but there are additional important topics that go beyond the most basic tasks an administrator must perform. Read on to learn about these.

Password Resets

The Add User functionality includes a *Require Password Reset* checkbox at the bottom of the Password form. The default password policy does not even allow administrators to deselect this option. As the administrator, however, you can modify the default password policy so that this box becomes usable.

1. Navigate to *Password Policies* in Control Panel → Users.
2. Click on the *Default Password Policy*.
3. Deselect the *Change Required* switcher in the Password Changes section. Now you can decide whether each user you add must reset his/her password.

See the article on Password Policies for more information on editing the default policy or creating your own.

Adding an Administrative User

If you're setting up a portal for the first time, you're likely to be using the default administrator account, the account of one of those famous Liferay Administrators, *Test Test* or her cousin, *Joe Bloggs*. Because these are default accounts, hackers know about them, so it's better to set up your own administrator account. Add a user with your information, then give your user account the same administrative rights as the default administrator's account:

1. Click the *Roles* link in the left navigation pane (in the *Edit User* page's *General* tab). This page of the form shows the Roles to which your account is currently assigned. No roles appear by default (the User role does not appear since it can't be removed).
2. Click *Select* under Regular Roles and assign the Administrator Role to your user account. A dialog box pops up with a list of all the regular (portal-scoped) Roles in the portal. Select the Administrator role from the list (click *Choose*). The dialog box disappears and the role is added to the list of Roles associated with your account. You are now a portal administrator. Log out and then log back in with your own user account.

Power Users: Users are not assigned the Power User role by default. The Power User role grants more permissions than the User role. If the User role is sufficient for you, ignore the Power User role. Alternatively, use it to provide a second level of User permissions and assign it to those Users. If there are certain custom permissions that you'd like all of your portal Users to have, you can grant these permissions to the User Role. You can also customize the default Roles a new User receives via *Default User Associations*. This is covered in the article on Instance Settings.

In production, you should always delete or disable the default administrator account to secure your portal.

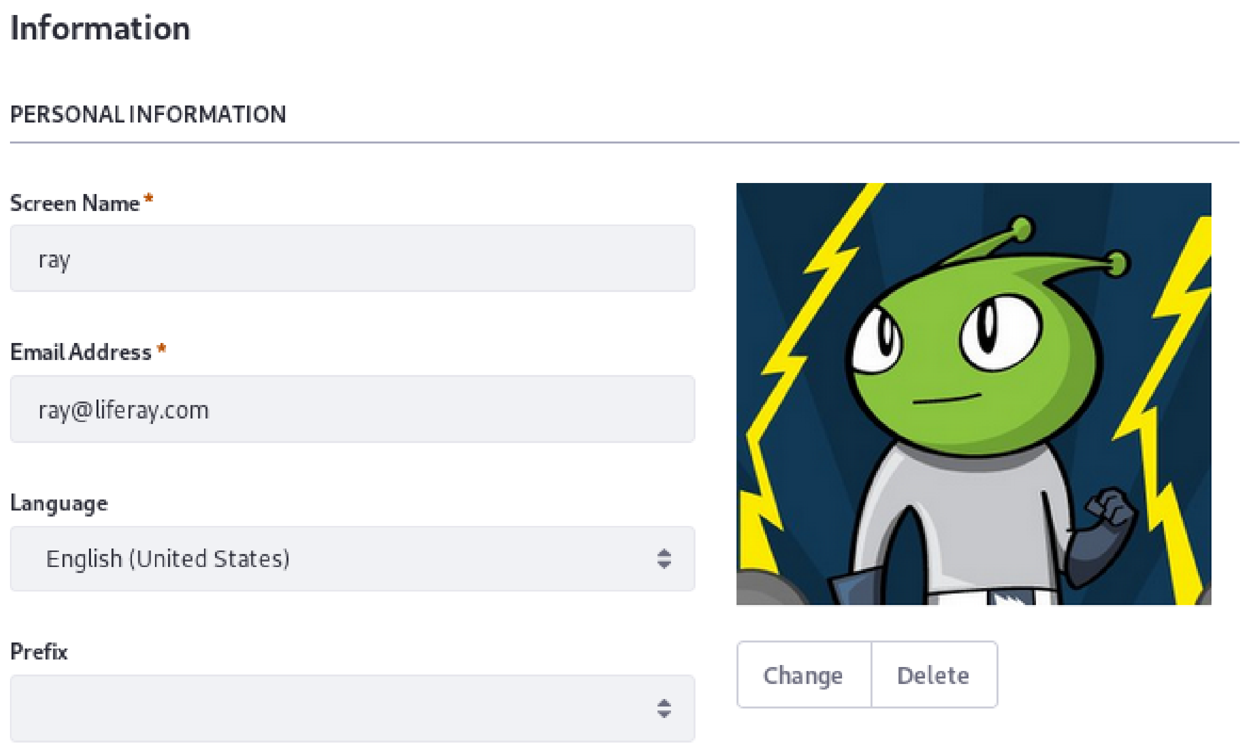

Figure 66.5: Upload images for user avatars in the Edit User form.

User Profile Pictures

Users have profile pictures. Administrative Users can upload images in the Edit User form, and any User can update her own account information, including image, from her personal site (*My account* → *Account Settings*).

If no image is explicitly uploaded for a User's profile picture, a default User icon is assigned as the User avatar. By default the User's initials are displayed (First Name then Last Name) over a random color.

Figure 66.6: The default user profile picture is an icon with the user initials over a randomly colored bubble.

If the initials-based approach for generating User profile pictures isn't suitable for your portal, disable the inclusion of Users' initials in the default icons:

1. Navigate to *Control Panel* → *Configuration* → *System Settings*.
2. In the Platform section, click *Users* → *User Images*.
3. Deselect *Use Initials for Default User Portrait*.

Figure 66.7: If you disable the default initials-based profile picture, this icon is used instead.

Now, instead of the default icon, which is a colorful circle containing the user's initials, the icon is a gray circle containing the approximate shape of a human being.

This is just the default. To override it with your own default image

1. Create at least one image that is a 100x100 px square. Place it somewhere on the application server's classpath.

2. Override the values of these portal properties:

```
image.default.user.female.portrait=com/liferay/portal/dependencies/user_female_portrait.png
image.default.user.male.portrait=com/liferay/portal/dependencies/user_male_portrait.png
```

3. Restart the application server.

To make the above work, you must place the new image on the classpath of the application server. For example, in Tomcat you could place images in the `tomcat/webapps/ROOT/WEB-INF/classes` folder and specify

```
image.default.user.female.portrait=user_female_portrait-override.png
image.default.user.male.portrait=user_male_portrait-override.png
```

in a `portal-ext.properties` file. Once you restart Tomcat, the new default images are used in place of the old ones.

Note: There's a way to adjust which initials are displayed and in what order, so you can make the default user icon (with the user initials) work for your locale. These settings are configured in a Language Settings module, so kidnap a friendly developer, give him a cup of coffee, and tell him the settings you want to change:

`lang.user.default.portrait=initials` sets the type of icon to use for avatars. The default value is *initials*. If set to initials, the next property configures which initials to display, and in what order. Alternatively, specify *image*, which gives you the same non-initials default image shown above.

`lang.user.initials.fields=first-name,last-name` determines which initials appear in the user portrait and in what order. The setting here only matters if `lang.user.default.portrait` is set to *initials*. Valid values are first name, middle name, last name, with first and last name as the defaults.

Numeric Screen Names

In prior versions, numeric user screen names were disabled out of the box via the default portal property

```
users.screen.name.allow.numeric=false
```

Other user management systems (LDAP, for example) did not have the same restriction, which made importing users more difficult. Administrators first had to set the above property to `true` before importing and hope that no screen names conflicted with site IDs. In 7.0, this property defaults to `true` and there's no danger of numeric screen names conflicting with site IDs:

```
users.screen.name.allow.numeric=true
```

This means you're free to set a user screen name to *24601*, or whatever other number you can think of, and imports from systems that allow numeric screen names go more smoothly. That's everything you need to know to take advantage of this feature. Keep reading to understand what enabled the change.

Because users have personal sites, the URL to user *24601*'s personal site is

```
http://localhost:8080/web/24601
```

Meanwhile, a default site URL to cleverly named *Test Site* is

```
http://localhost:8080/web/test-site
```

There's no conflict here, but two conditions could easily lead to one:

1. *Test Site*'s group ID matches the number chosen for the user's screen name. Each site has a unique numeric identifier in the database, called group ID. There's nothing stopping it from matching the user's numeric screen name, so it could easily be `24601` just like the hypothetical user above.

2. A site administrator comes along and changes the site's friendly URL to match its `groupId`. Hello, URL conflict! Now the site's URL matches the user's URL:

   ```
   http://localhost:8080/web/24601
   ```

This conflict is no longer possible. In 7.0, a site's friendly URL is not allowed to be numeric. See for yourself:

1. Navigate to the site's *Configuration* → *Site Settings* → *Site URL* section.

2. In the Friendly URL section, enter *24601* and save the form. A failure message appears if you don't have a user with the matching screen name:

   ```
   The friendly URL may conflict with another page.
   ```

   You'll see this failure message if there's an existing conflict with a user screen name:

   ```
   Please enter a unique friendly URL. Site [user-first-name user-last-name] has the same friendly URL.
   ```

Next, learn about collecting users in organizations.

66.4 Organizations

An *Organization* groups *Users* hierarchically. For example, you can model a company's departments (i.e., Human Resources and Customer Support) with organizations. Organizations often have their own sites. The *how-to* portion of managing Organizations is in the next article, Managing Organizations. This article contains important conceptual information on what Organizations are and when they're needed.

Many simple portal designs don't use Organizations at all; they only use sites. The main purpose of Organizations is to enable distributed User management. Portal administrators can delegate some user management responsibilities to Organization administrators. If you don't anticipate needing to delegate User management responsibilities, your portal design probably doesn't need to include Organizations.

User Groups and Organizations: It's easy to confuse User Groups (covered in a separate article) with Organizations since they both group Users. User Groups are an ad hoc collection of Users, organized for a specific function. In the Lunar Resort, if you wanted a group of bloggers, for example, it wouldn't make sense to assign the Sales Department the role of blogging (see the article on Roles if you're not sure what they are). The Sales Department users could blog whenever a new T-shirt design became available in the Lunar Resort store, but they probably wouldn't be as diligent about announcing the new Rover Racing schedule. Instead, creating a User Group containing one individual from each department who is responsible for blogging would make more sense. Read the article on User Groups to learn more about how to use them in your portal.

When to Use Organizations

To decide whether your portal design should include Organizations, think about its function. A photo-sharing web site could be powered by Sites only. On the other hand, Organizations are useful for corporations or educational institutions since their users can be placed into a hierarchical structure. Don't think that Organizations are only for large enterprises, though. Any group hierarchy, from large government agencies all the way down to small clubs, can be modeled with Organizations. Also, don't think that you must decide between an Organization-based structure or a Site-based structure for assembling your portal's Users. Users can belong both to Organizations and to independent Sites. For example, a corporation or educational institution could create a social networking site open to all Users, even ones from separate Organizations.

To illustrate what an Organization is, consider a potential Organization of the Lunar Resort's Intranet. The company hierarchy has three tiers: The Lunar Resort, its departments, and divisions within each department.

- Lunar Resort–The top-level Organization.
  - Physical Plant Department–Department of users that keep the place running.
    - Grounds Crew–Users that maintain the grounds.
    - Janitorial Crew–Users who keep the resort clean.
    - Mechanical Crew–Users who fix stuff, like lunar rovers.
  - Recreation Department–A department that makes sure much fun is had by guests of the Lunar Resort.

    * Golf Instructors–Teach guests how to golf on the moon.
    * Rover Race Instructors–Teach guests how to drive the lunar rovers.
    * Lunar Sherpas–Lead guests on moon hikes.
  - Sales Department–A department of users who sell things to Lunar Resort guests.
    * Up-sale Group–Make sure guests know how easy it is to improve their stay by spending more money.
    * Souvenir and Memorabilia Group–Peddle souvenirs to Lunar Resort guests.
    * Retail Group–Maintain the Lunar Resort store, which contains basic necessities, since guests are coming all the way from Earth.
  - Sentient Organism Resources Department–Department of Users that hire, fire and regulate intra-company relationships. We'd call it Human Resources, but what's stopping Martians from applying? Nothing!

Each department is a sub-Organization of the resort, and each division is a sub-Organization of the department.

What can Organization Administrators Do?

Whenever you have a collection of Users that fits into a hierarchical structure, you can use Organizations to model those Users. Organization administrators can manage all the Users in their Organization *and* in any sub-Organization. Referring to the hierarchy above, for example, an Organization administrator of the Lunar Resort could manage any Users belonging to the resort itself, to any of the departments, or to any of a department's subdivisions. An Organization Administrator of the Physical Plant Department can manage any Users belonging to the Physical Plant Department itself, or to the Grounds Crew, the Janitorial Crew, or the Mechanical Crew. However, an administrator of the Physical Plant Department can't manage Users belonging to the Recreation Department or the Retail Group organization.

Organizations and sub-Organization hierarchies can nest to unlimited levels. Users can be members of one or many Organizations. The rights of an Organization administrator apply both to his/her Organization and to any child Organizations. Members of child Organizations are implicit members of their parent Organizations. This means, for example, that members of child Organizations can access the private pages of their parent Organizations. This behavior can be customized in the `Organizations` section of the portal-ext.properties file where the properties specific to organizations are listed.

Since Organizations are designed for distributed user administration, Organization Administrators have an entirely different set of privileges than Site Administrators. Site Administrators are responsible for the pages, portlets, content, and membership of their Sites. To this end, they can set the membership type to Open, Restricted, or Private. They can also add Users to or remove Users from their Sites but cannot manage the Users themselves. If an Organization has a Site attached to it, the Organization Administrator has the same rights as a Site Administrator for managing the Site's content, but an Organization Site's members are the members of the Organization. Thus Organization administrators have more user management permissions than Site administrators: they can edit users belonging to their Organization or any sub-Organization. They cannot add existing portal Users to their organization, but they can create new Users within their Organization. Only portal administrators can add existing users to an Organization.

Organization Administrators can't access the Control Panel by default, but it's not necessary. In their personal Sites, Organization administrators can click the *My Organizations* link to gain access to any Organizations they manage.

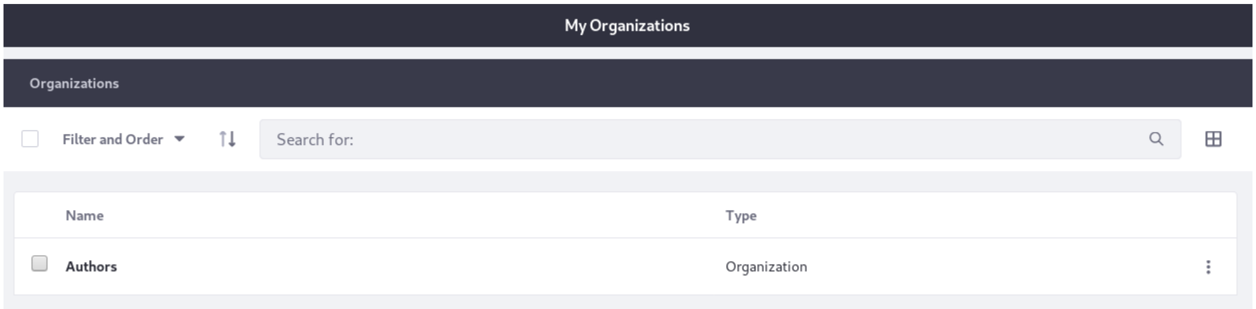

Figure 66.8: The My Organizations application lets Organization Administrators manage their organizations in their personal site.

Organization Roles and Permissions

A huge time-saving benefit of including Organizations into your portal design is that Organization administrators can assign Organization-scoped Roles to members of the entire Organization. For example, consider an IT Security group in a corporate setting. You could have a sub-Organization of your IT organization that handles security for all applications company-wide. If you grant the IT Security Organization the portal administrator Role, all members of the Organization get administrative access to the entire system. Suppose further that a User in this Organization was later hired by the Human Resources department. The simple act of removing the User from the IT Security Organization also removes the User's administrative privileges, since the privilege came from the IT Security Organization's Role. By adding the User to the HR Organization, any roles the HR Organization has (such as access to a benefits system in the portal) are transferred to the User. In this manner, you can design your portal to correspond with your existing organization chart and Users' permissions are granted according to their positions in the chart.

Of course, this is only one way to set things up. If you have more complex requirements for permissions within an Organization, you can create custom Organization-scoped Roles to assemble the permissions you wish to grant to particular Users. Alternatively, attach a Site to your Organization and use Site Teams to assemble the sets of permissions (see below). See the Roles and Permissions article for more detail.

Organization Sites

Does your Organization need to have its own Site? If an organization has an attached Site, the Organization's administrators are treated as the Site administrators. This means that they can manage the pages, portlets, and content of the Site as well as the Users of the Organization. Members of an Organization with an attached Site are treated as members of the Organization's Site. This means that they can access the private pages of the Organization's Site, along with any portlets or content there. Attaching Sites to Organizations allows portal administrators to use Organizations to facilitate distributed portal administration, not just distributed User administration.

That's a lot of information on Organizations. Next, learn how to create and manage Users and Organizations.

66.5 Managing Organizations

If you're not entirely sure what Organizations are or whether you need them in your portal, start here. This article gets right to the practical stuff: how to manage Organizations.

Adding Organizations

Add an Organization (perhaps start by adding the *Physical Plant Department* organization to the Lunar Resort):

1. Click *Users and Organizations* from Control Panel → Users.
2. Go to the *Organizations* tab and click the *Add* button. Fill out the Name field at a minimum.
3. To attach a Site when you create an Organization, click on *Organization Site* and check the *Create Site* box. If you don't know right now if your organization needs a site, that's fine. You can always add one later.

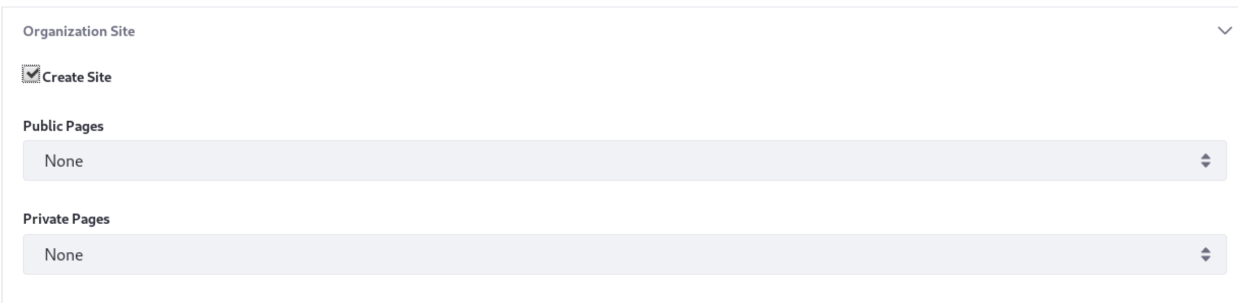

Figure 66.9: Fill out the add Organization form and click *Save*.

4. If you're creating a child Organization, use the Parent Organization *Select* button to select an Organization in the system to be the direct parent. Click the *Remove* button to remove the currently configured parent.
5. Click *Save* when finished filling out the Add Organization form.

As when creating a new user, once you submit the form a success message appears and you have access to a new form which lets you enter additional information about the Organization. Organizations can have associated multiple email addresses, postal addresses, web sites, and phone numbers. The Services link can be used to indicate the operating hours of the Organization, if any.

Tip: After creating an Organization, assign the desired user to the Organization Owner Role. The Organization Owner can do everything that an organization Administrator can. In addition to their full administrative rights within the Organization, they can do these things:

- Appoint other Users to be Organization Administrators
- Appoint other Users to be Organization Owners
- Remove the memberships of other Organization Administrators or Owners

Organization Administrators can't make these Role assignments and can't manage the memberships of other Organization Administrators or Owners.

Editing Organizations

To edit an Organization, go to the Users and Organizations section of the Control Panel and click the *Organizations* tab. All active Organizations are listed. Click the *Actions* button next to an Organization. This shows a list of actions you can perform on this Organization.

- *Edit* lets you specify details about the Organization, including addresses, phone numbers, email addresses and websites.
- *Manage Site* lets you create and manage the public and private pages of the Organization's site. This only appears for Organizations that have attached Sites.
- *Assign Organization Roles* lets you assign Organization-scoped Roles to Users. By default, Organizations are created with three Roles: Organization Administrator, Organization User and Organization Owner. You can assign one or more of these Roles to Users in the Organization. All members of the Organization automatically get the Organization User Role so this Role is hidden when you click Assign Organization Roles.
- *Assign Users* lets you search and select Users to be assigned to this Organization as members.
- *Add User* adds a new User and assigns the User as a member of this Organization.
- *Add Organization* lets you add a child Organization to this Organization. This is how you create hierarchies of Organizations with parent-child relationships.
- *Delete* removes this Organization. Make sure the Organization has no Users in it first. You'll be prompted for confirmation that you want to delete the Organization. If there are Users in the Organization or if there are sub-Organizations, you must remove the Users and delete the sub-Organizations before deleting the parent Organization.

If you click the Organization name you can view both a list of Users who are members of this Organization and a list of all the sub-Organizations of this Organization.

Organization Types

By default, Liferay DXP only includes the *Organization* type. Configure the existing type or add additional types using the aptly named Organization Type entry in System Settings. There are two main reasons to configure Organization types:

1. Organizations usually correlate to real-life hierarchical structures. Calling them by their real names is helpful for administrators and Users. In the Major League Baseball (MLB) example, *League*, *Division*, and *Team* Organization types are useful.
2. Enforce control over which Organizations can be top level Organizations and the type of sub-Organization allowed for each parent Organization type. For example, MLB would not allow Division Organization types to be sub-Organizations of Team Organizations.

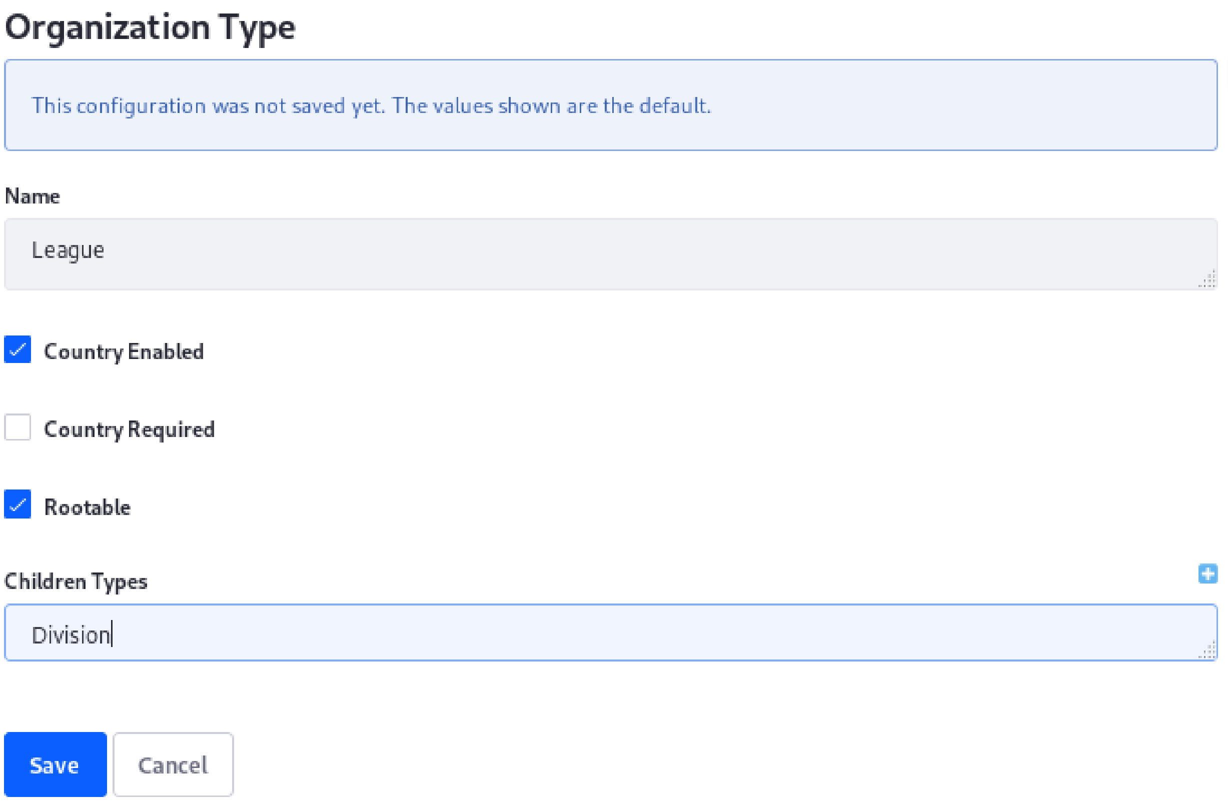

Figure 66.10: Create new organization types through the System Settings entry called Organization Types.

Check out the configuration options that configure the default *Organization* type and then configure an additional type.

To add another Organization type called *League*, enter these options into the configuration form:

Name: ***League*** Adds League to the list of Organization types that appear in the Add Organization menu.

Country Enabled: ***True*** Enables the Country selection list field on the form for adding and editing League types.

Country Required: ***False*** Specifies that the *Country* field is not required when adding a League.

Rootable: ***True*** Enables Leagues as a top level Organization. Limit League to sub-Organization status by excluding this property.

Children Types: ***Division*** Specifies Division as the only allowable sub-Organization type for the League parent type.

Once you configure additional Organization types and click Save, you'll find your new type(s) available for selection in the Add Organization form.

Users can join or be assigned to Sites when they share a common interest. Users can be assigned to Organizations when they fit into a hierarchical structure. User groups provide a more ad hoc way to group users than sites and Organizations. You'll look at them next.

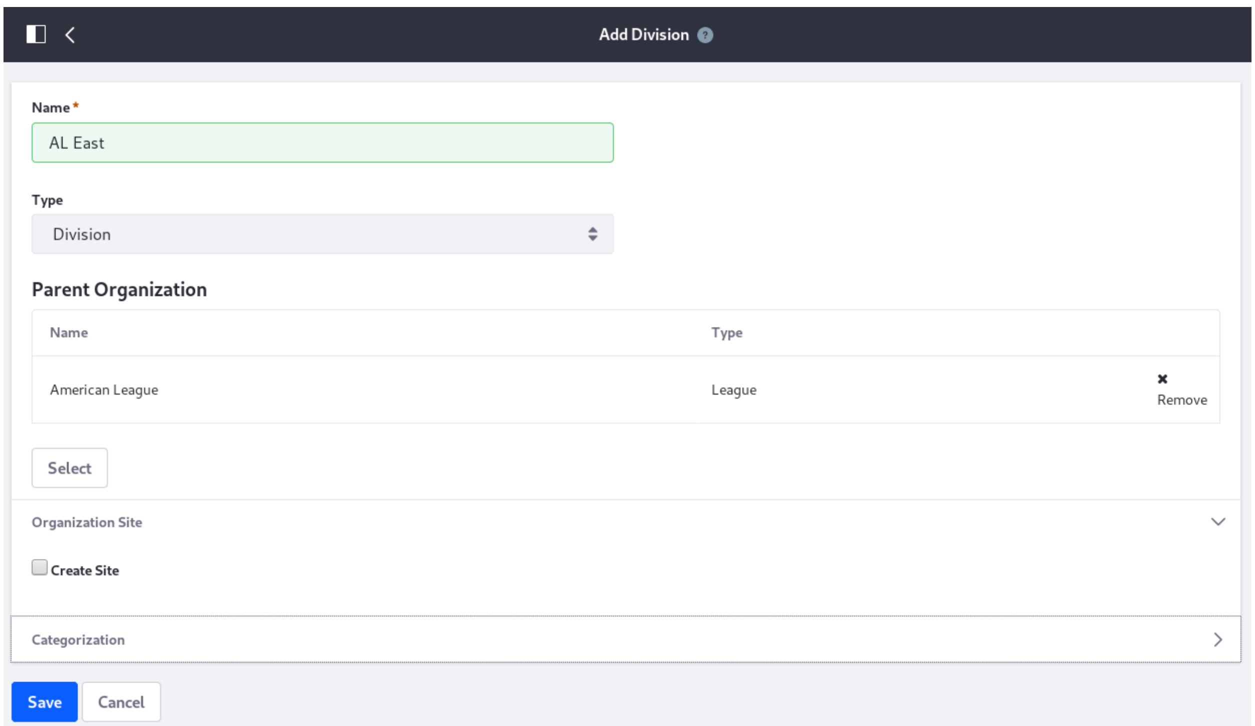

Figure 66.11: Custom configuration types are available in the Add Organization form.

Chapter 67

Roles and Permissions

If a *Role* were to win a Grammy or an Oscar or some other ego-feeding popularity contest, it better remember to thank all its *permissions* groupies during the acceptance speech, because they're the ones doing the real work. The role is just the pretty face, so to speak.

Roles collect permissions that define a particular function, according to a particular scope. Roles collect permissions, and Users are assigned to Roles.

Note: Roles are assigned to Users, but it's tedious to assign each User to role intended for lots of Users. Recall that Users are grouped in Sites, Organizations, and User Groups. Implicitly assign regular scoped permissions to Users by assigning a role directly to one of these user groupings.

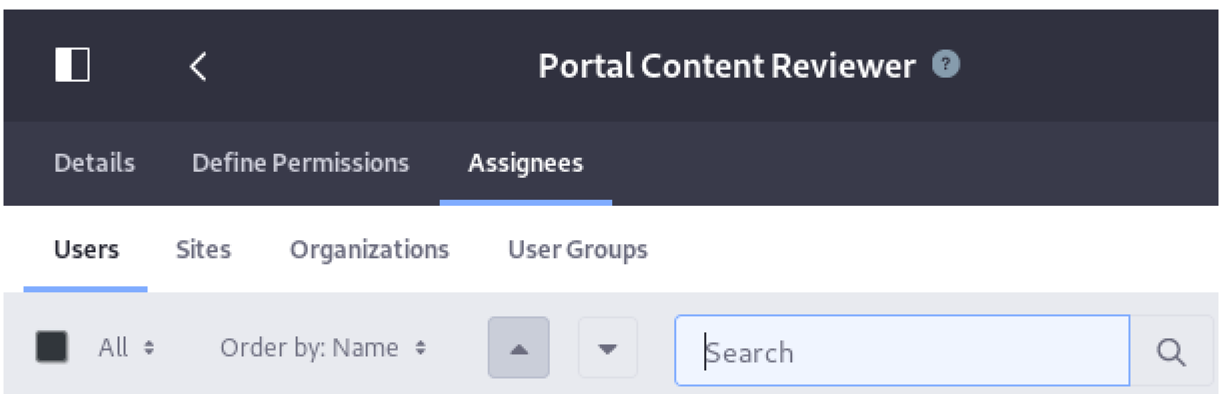

Figure 67.1: Assign Users to a role, directly or by their association with a Site, Organization, or User Group.

Take a Message Board Administrator Role, for example. A Role with that name should have permissions relevant to the specific Message Board portlets delegated to it. Users with this Role inherit the permissions collected underneath the umbrella of the Role.

In addition to regular Roles, Site Roles, and Organization Roles, there are also Teams. Teams can be created by site administrators within a specific Site. The permissions granted to a Team are defined and applied only within the Team's site. The permissions defined by regular, Site, and Organization Roles, by contrast, are defined at the global level, although they are applied to different scopes.

Regular role Permissions are defined at the global level and are applied at the global scope.
Site role Permissions are defined at the global level and are applied to one specific Site.
Organization role Permissions are defined at the global level and are applied to one specific Organization.
Team Permissions are defined within a specific Site and are assigned within that specific Site.

Note: Some permissions cannot be handled from the control panel. Asset-level permissions (for instance, permission to edit an individual blog post, or view a folder in the Documents and Media library) are managed from the individual asset. See Widget Permissions for details.

67.1 Deleting Asset Containers

A Web Content Folder contains Web Content articles. The Web Content Folder is an asset container, and the Web Content Article is an asset. It's possible to give a Role permission to delete an asset container without giving the Role permission to delete individual assets. In that case, beware: if a Role assignee deletes an asset container with individual assets in it, the individual assets themselves are deleted as well.

Besides Web Content Folders, examples of asset containers include Bookmarks Folders, Message Boards Categories, Wiki Nodes, and Documents and Media Folders.

You might not need to create a Role for a certain functionality. Liferay provides many pre-configured Roles for your convenience.

67.2 Default Liferay Roles

In the Roles Application appears a list of all the Roles in Liferay, by scope.

These are some of the pre-configured regular Roles:

- Guest: The Guest role is assigned to unauthenticated users and grants the lowest-level permissions.
- User: The User role is assigned to authenticated Users and grants basic permissions (mostly *Add to Page* permissions for their own Sites).
- Power User: The Power User Role grants more permissions than the User Role. It's an extension point for distinguishing regular Users from more privileged Users. For example, you can set things up so that only Power Users have personal sites.
- Administrator: The administrator Role grants permission manage the entire portal, including global portal settings and individual Sites, Organizations, and Users.

These are some of the pre-configured site roles:

- Site Member: The Site Member Role grants basic privileges within a Site, such as permission to visit the Site's private pages.

- Site Administrator: The Site Administrator Role grants permission to manage *almost* all aspects of a Site including site content, site memberships, and site settings. Site Administrators cannot delete the membership of or remove roles from other Site Administrators or Site Owners. They also *cannot* assign other Users as Site Administrators or Site Owners.
- Site Owner: The Site Owner Role is the same as the Site Administrator Role except that it grants permission to manage *all* aspects of a Site, including permission to delete the membership of or remove Roles from Site Administrators or other Site Owners. They *can* assign other users as Site Administrators or Site Owners.

These are some of the pre-configured organization roles:

- Organization User: The Organization User role grants basic privileges within an Organization. If the Organization has an attached Site, the Organization User Role implicitly grants the Site member Role within the attached Site.
- Organization Administrator: The Organization Administrator Role grants permission to manage *almost* all aspects of an Organization including the Organization's Users and the Organization's Site (if it exists). Organization Administrators cannot delete the membership of or remove Roles from other Organization Administrators or Organization Owners. They also *cannot* assign other Users as Organization Administrators or Organization Owners.
- Organization Owner: The Organization Owner Role is the same as the Organization Administrator Role except that it grants permission to manage *all* aspects of an Organization, including permission to delete the membership of or remove Roles from Organization Administrators or other Organization Owners. They *can* assign other Users as Organization Administrators or Organization Owners.

Tip: It's easy to overlook the differences between owner type roles and administrator type roles for Sites and Organizations. Site and Organization administrators *cannot* remove the administrator or owner Role from any other administrator or owner, and they *cannot* appoint other Users as site or organization administrators or owners.

In contrast, site and organization owners *can* do those things.

Roles, and the permissions granted with their assignment, are foundational components in Liferay. Understanding their uses and configuration enhances your ability to configure Liferay DXP to suit your organizational needs.

67.3 Managing Roles

You manage Roles and Permissions in the Control Panel (*Control Panel → Users → Roles*). There you'll find an application for creating Roles, granting them permissions, and assigning Users to them. Roles can be scoped by portal, Site, or Organization.

Defining a Role's permissions is a topic deserving its own article. Read here about defining a role's permissions.

Creating Roles

Determine the scope of the Role you must create. Roles can be scoped globally (Regular Roles), to a specific Site (Site Roles), or to an Organization (Organization Roles).

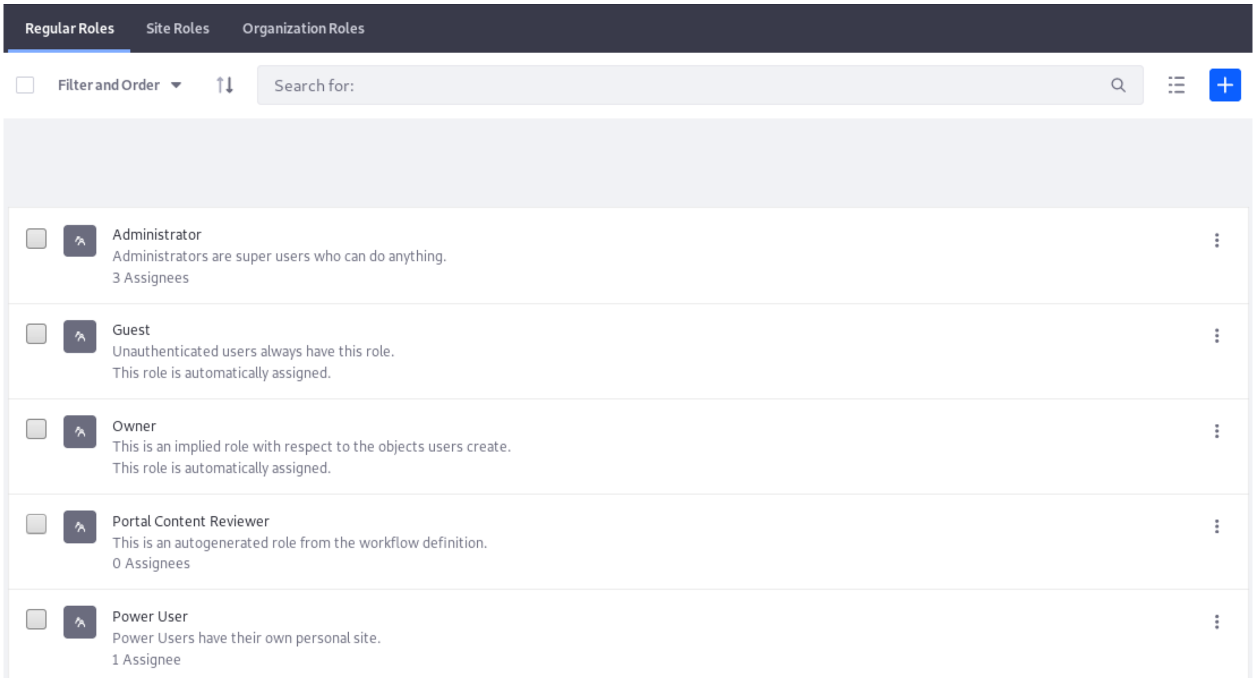

Figure 67.2: The Roles application lets you add and manage roles for the global (Regular), Site, or Organization scope.

To create a Role:

1. Click the tab for the proper Role scope, then click the *Add* () button.
2. Enter a title and description. The title field is required but the description is optional.
3. Enter a Key, if desired. This required field provides a key that can be used to refer to the Role programmatically. It's auto-populated with the title text, but you can override it if desired.
4. Click *Save*.

Now the Role is present in the database and ready for further configuration.

Assigning Users to a Role

Assign users to a Role in the Assignees tab of the Add/Edit Role form. Roles are assigned to Users, Sites, Organizations, or User Groups. Here's how to assign the User Group Manager Role created in the last section to Users:

1. In the Assignees tab of the Add/Edit Role form, click the second level tab for *Users*.
2. Click the Add button ().
3. Select the Users you want to add to the Role and click *Add*.

If assigning a group, note that all Users assigned to that group inherit the Role as well.

That's a good start, but your Role isn't worth the database row it occupies without defining permissions for the Role. Read the next article to learn how.

67.4 Defining Role Permissions

Roles collect permissions, so when Users are given a Role, they receive all the permissions defined by the Role.

If you create a Role with permission to access something in the Control Panel, keep in mind that the *View Control Panel Menu* permission is automatically granted.

Consider a Role called User Group Manager. Define the permissions for the User Group Manager Role so that assigned Users can add Users to or remove Users from any User Group:

1. Go to the Control Panel and then click on *Users* → *Roles*.
2. On the Regular Roles screen, click *Add* ( ).
3. After naming your Role, click *Save*.
4. Click on the *Define Permissions* tab.
5. Drill down in the menu on the left to *Control Panel* → *Users* → *User Groups*.
6. Under the *General Permissions* heading, flag *Access in Control Panel* and *View*. This lets user group managers access the User Groups Control Panel portlet and view existing User Groups.
7. Since you want User Group managers to be able to view User Groups and assign members to them, also check the *Assign Members* and *View* permissions under the *Resource Permissions* → *User Group* heading.
8. There's one last necessary permission you might not think of in association with this Role. In *Control Panel* → *Users* → *Users and Organizations*, User Group managers need *View* permission on the User resource. Grant this permission.
9. Click *Save*.

Now the User Group Manager Role has all the permissions necessary for adding Users to User Groups. After all, User Group managers can view User Groups, assign members, and access User Groups in the Control Panel. The permission to view Users in the Control Panel was necessary because you must view Users to assign them as members of a Role. Without this permission, User Group managers see an empty list if they try to add Users to a Role.

Note: The Roles application in the Control Panel is not the only place where permissions are configured. You can configure a Role's permissions on a resource at a more granular level. For example, from a particular application instance, click its *Options* ( ) menu and select *Permissions*. There you can configure permissions for the resource that overlap with those configured in the Control Panel's Roles application. However, permissions granted or removed in the Control Panel override those made at the more granular level.

There are three categories of permissions: *Control Panel*, *Site Administration*, and *User*. By default, Users can manage their User accounts via the permissions belonging to the User category. Site Administrators can access the site administration tools belonging to the Site Administration category. Portal Administrators can access the entire Control Panel. For custom Roles, you can mix and match permissions from as many categories as you like.

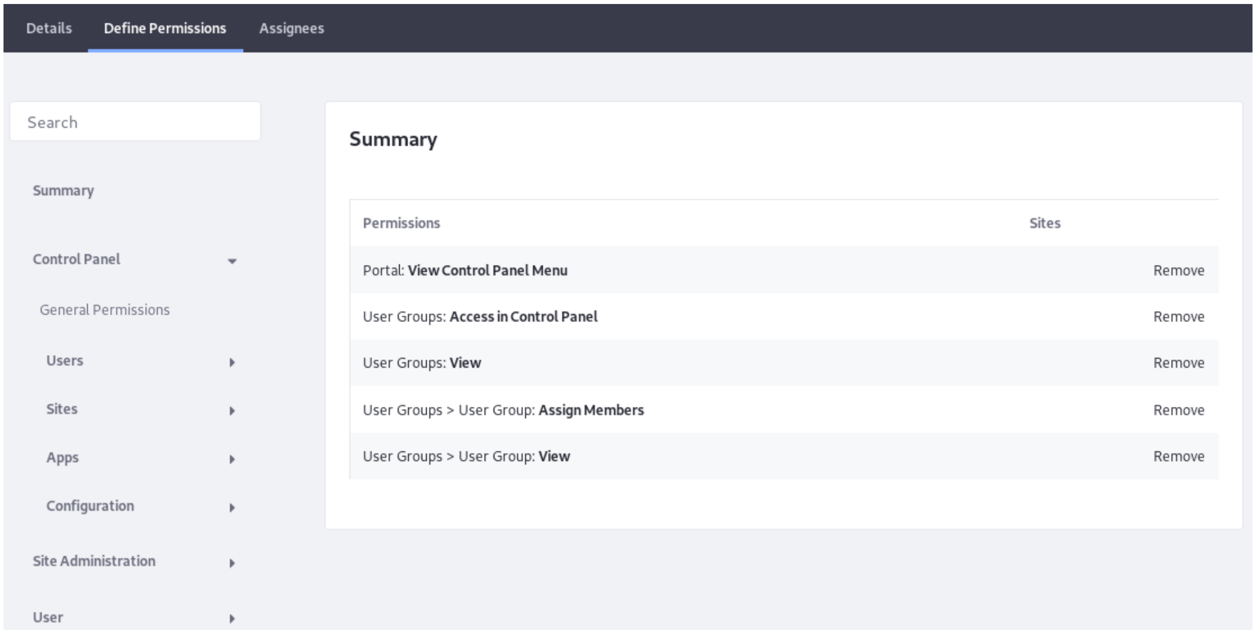

Figure 67.3: When defining permissions on a role, the Summary view provides a list of permissions that have already been defined for the role. The area on the left side of the screen lets you drill down through various categories of permissions.

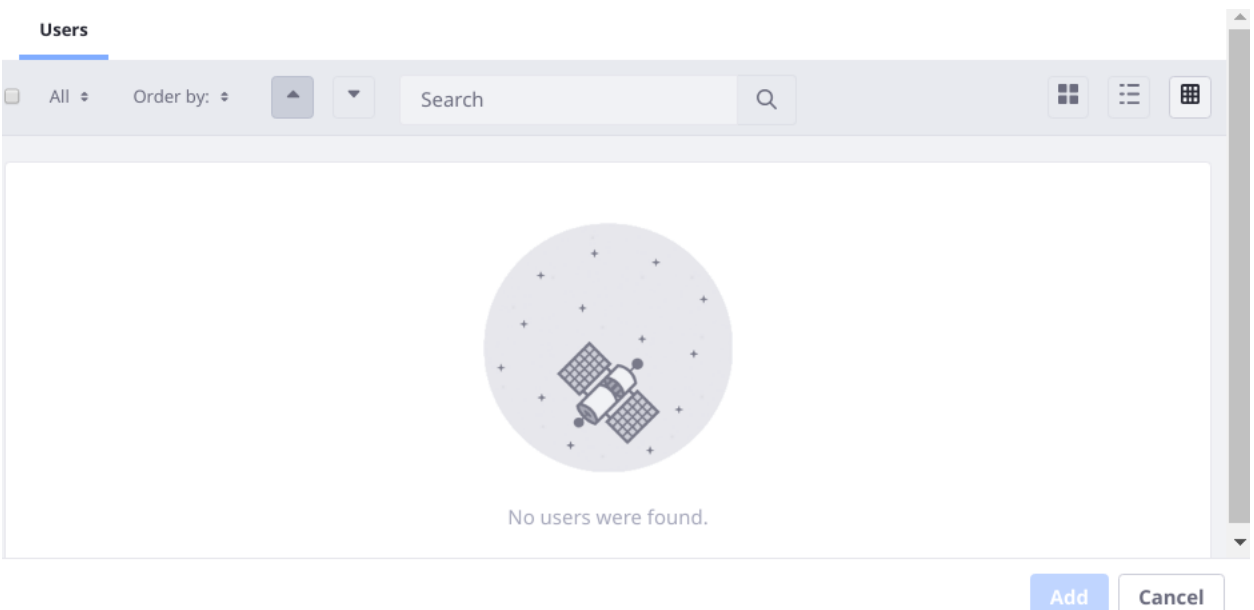

Figure 67.4: Users assigned to the User Group Manager role can't find any users to add unless they have view permissions on the User resource.

The permissions in the Site Administration → Applications categories govern the content that can be created by portlets such as the Wiki and Message Boards. If you pick one of the portlets from this list, you see options for defining permissions on its content. For example, if you pick Message Boards, you see permissions for creating categories and threads or deleting and moving topics.

Site application permissions affect the application as a whole. Using the Message Boards as an example, an application permission might define who can add the Message Boards portlet to a page.

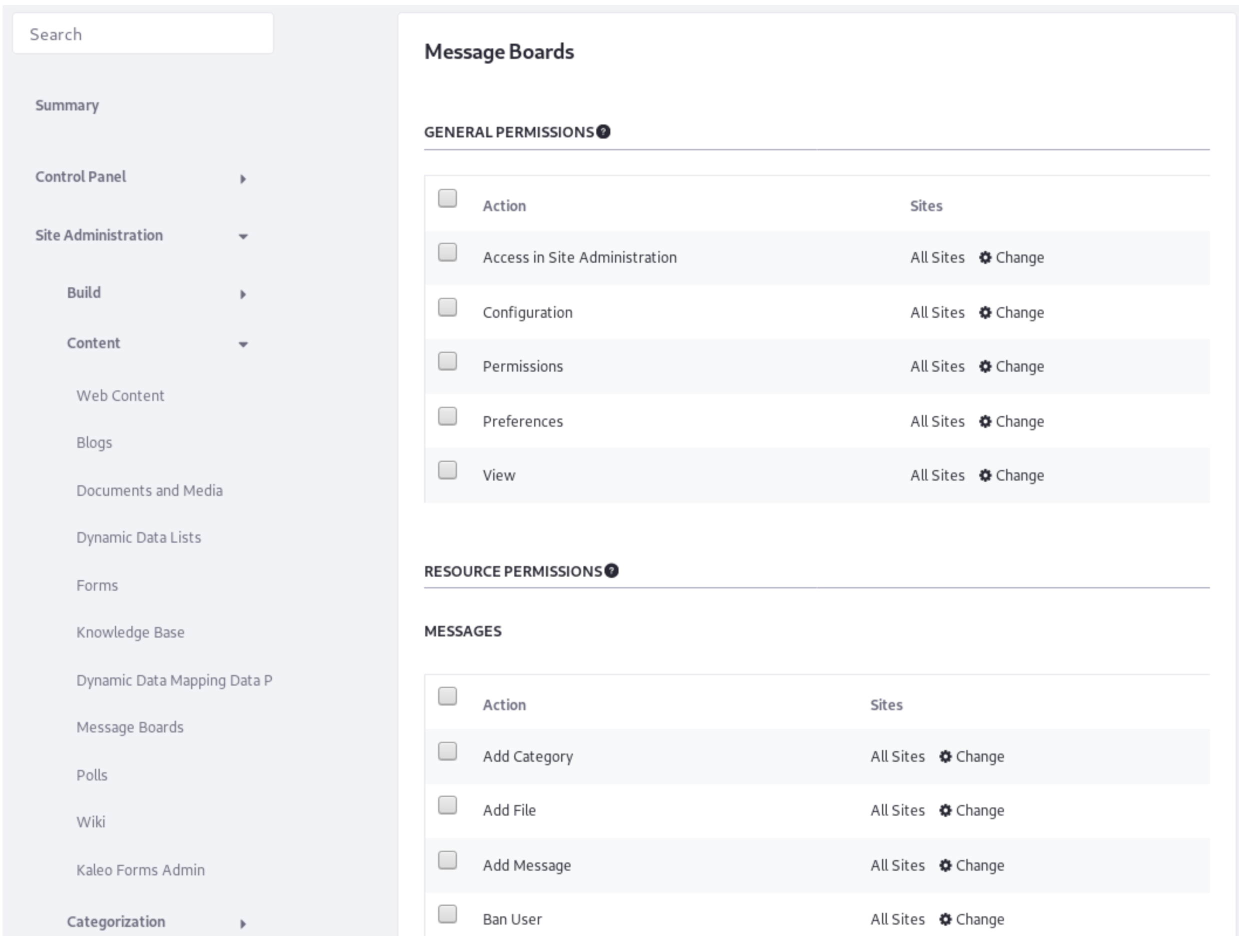

Figure 67.5: You can fine-tune which actions are defined for a role within a specific application like the Message Boards.

The Control Panel permissions affect how the Control Panel appears to the User in the Control Panel. The Control Panel appears differently for different Users, depending on their permissions. Some Control Panel portlets have a Configuration button, and you can define who gets to see that. You can also fine-tune who gets to see various applications in the Control Panel.

If you want to change the scope of a permission, click the *Change* link next to the gear icon next to the permission and then choose a new scope. After you click *Save*, you'll see a list of all permissions currently granted to the Role. From the Summary view, you can add more permissions or go back to the Role Application default view by clicking on the *Back* ( ) icon.

Sometimes you might find that a certain permission grants more or less access than what you expected—always test your permissions configurations!

Delegating Social Activities Configuration

There's a permission that allows Site administrators to delegate responsibility for configuring social activities to other Users. To add this permission to a Role, click *Actions* next to the desired Role and select *Define Permissions*. Find the *Site Administration* → *Configuration* → *Social Activity* permissions category. Flag all of the permissions and then click *Save*:

- Access in Site Administration
- Configuration
- Permissions
- Preferences
- View

Once these permissions are granted, Role assignees can manage the site's Social Activities.

Roles allow portal administrators to define various permissions in whatever combinations they like. This gives you as much flexibility as possible to build the Site you have designed.

CHAPTER 68

MANAGING USER DATA

Internet users are increasingly and justifiably concerned about how their personal data is processed by the systems they use. The enforcement of GDPR is a crystallization of these concerns into legislative action. Companies processing the personal data of EU residents must adopt appropriate measures to protect User data. Of course, legal requirements like those in GDPR only explain one reason for companies to develop policies for ensuring their users' right to privacy. The market demands site owners show higher levels of responsiveness to User inquiries into how their data is stored and processed.

Liferay is aware of the need for functionality to address User data management, and added two important features toward this end:

Erase and/or anonymize data associated with a User Administrative Users go through a step by step process, choosing to erase certain pieces of data and anonymize others.

Export a User's personal data Export ZIP files containing the data associated with a User.

These features are tools that get you closer to meeting two of GDPR's technically challenging requirements, the *right to data portability* and the *right to be forgotten*.

Note: It is Liferay's sincerest hope that through the User Management functionality of Liferay DXP, companies processing the personal data of their website's users can satisfy the requirements of GDPR. However, the tools discussed here and anywhere else in the documentation, including those directly aimed at addressing GDPR requirements (as in this article) do not guarantee compliance with the legal requirements of GDPR. Each company or individual whose website processes user personal data and is under the jurisdiction of GDPR must carefully determine the precise steps necessary to ensure they are fully compliant with GDPR.

68.1 Anonymizing Data

Deleting User data is the safest way to honor *right to be forgotten* requests. When User data must be preserved, automatic anonymization of the data is in order. Users being anonymized must have their identifiers (for example, User ID and User Name) removed from content they've interacted with. However, portal content usually requires this information for its applications to work properly. Therefore, the User's identifiers must be replaced by something, or someone. Meet the new User, *Anonymous Anonymous*, identity swapper *extraordinaire*. This deactivated User is dedicated to be the User whose identifiers are assigned to anonymized content. This identity swap is an important step in the anonymization process, but additional manual intervention may be necessary to truly achieve anonymization.

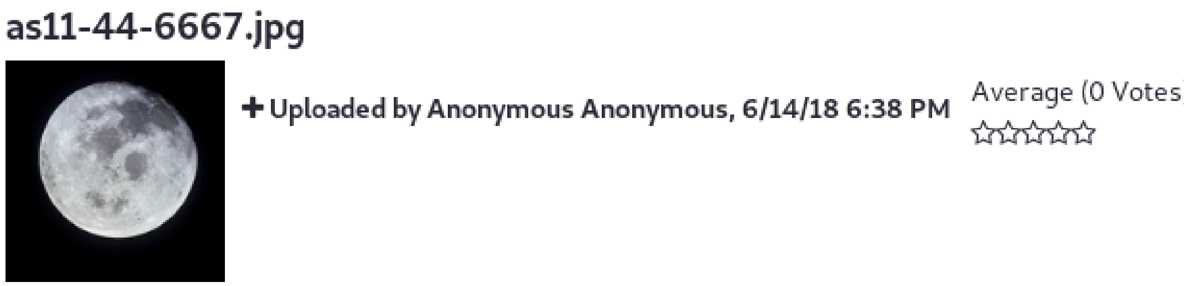

Figure 68.1: Anonymized content is presented with the User Anonymous Anonymous's identifying information.

Here at Liferay, we've grown fond of *Anonymous Anonymous*. If you'd rather start from scratch or assign an existing User to be the Anonymous User, get rid of *Anonymous Anonymous* and configure your own Anonymous User.

The anonymous user is programmatically created for each instance the first time an Administrator clicks *Delete Personal Data* from a User's Actions menu (⋮). If you haven't yet done that, no Anonymous User exists.

The easiest way to set up a new User as the Anonymous User is to edit an existing Anonymous User configuration, passing in the new Anonymous User's User ID.

To edit an existing configuration:

1. Go to Control Panel → Configuration → System Settings → Users → Anonymous User.

2. Edit the existing configuration, providing a different User ID.

   Get the User ID from Control Panel → Users → Users and Organizations. Click on the User and find the User ID in the Information screen of the Edit User application.

3. Click *Update*.

To create a new Anonymous User:

1. Create a User use for data anonymization. Alternatively, you can use an existing User.

2. If there's already an Anonymous User configured for the instance, there are two ways to remove it:

   Delete the User entirely. Deleting the User simultaneously deletes its configuration as the Anonymous User. Go to Control Panel → Users → Users and Organizations. If it's an active User, first deactivate, then delete the User. The default Anonymous Anonymous User is deactivated by default. Simply delete the User in this case. Click the Actions button (⋮) and select *Delete*.

   If you don't want to delete the User, just delete the User's configuration as the Anonymous User. Go to Control Panel → Configuration → System Settings → Users → Anonymous Users.

3. Add a new Anonymous User configuration. Click the *Add* button.

4. Fill out the two fields, Company ID and User ID.

   Get the Company ID from Control Panel → Configuration → Virtual Instances. The Instance ID and Company ID are the same.

   Get the User ID from Control Panel → Users → Users and Organizations. Click on the User and find the User ID in the Information screen of the Edit User application.

There can only be one Anonymous User configured for each instance.

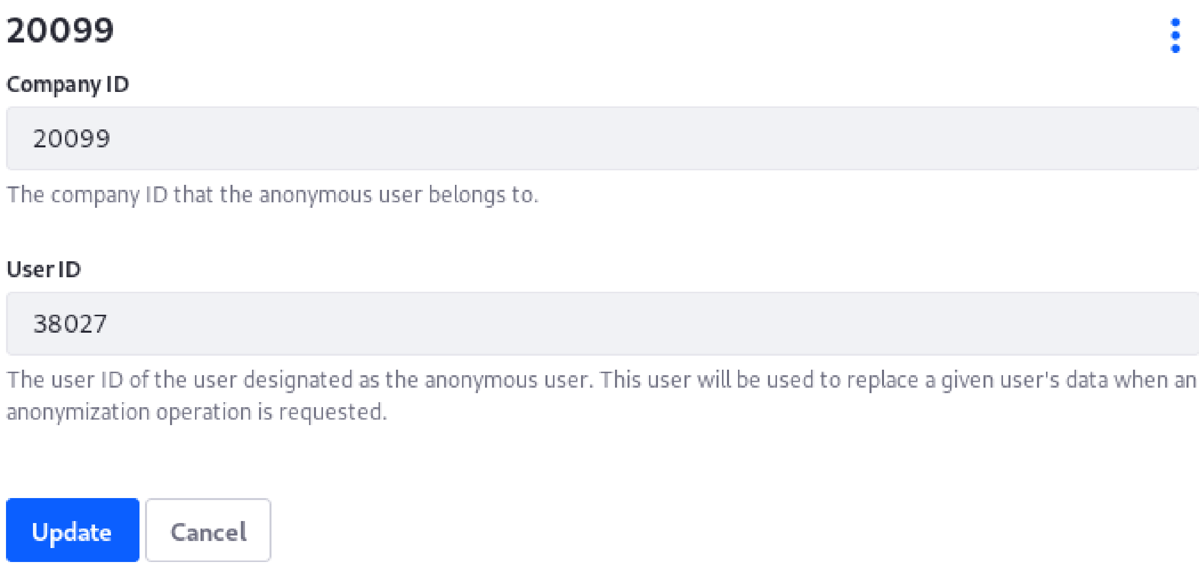

Figure 68.2: Assign your own Anonymous User from Control Panel → Configuration → System Settings → Users → Anonymous User.

68.2 Manual Anonymization

Anonymizing just the User's identification fields is often not enough. If a User named Ziltoid Omniscient complains about The Lunar Resort's coffee in a Message Boards Message and in it signs the post with *Supreme Leader of Ziltoidea 9*, anonymizing this post would remove the User's name (Ziltoid Omnisicent) and replace it with Anonymous Anonymous, but searching the internet for *Ziltoidea 9* quickly reveals that the post was written by Ziltoid the Omniscient. There can be user-entered personal data within the content of an application. You must manually edit such content to remove identifying details.

68.3 Sanitizing User Data

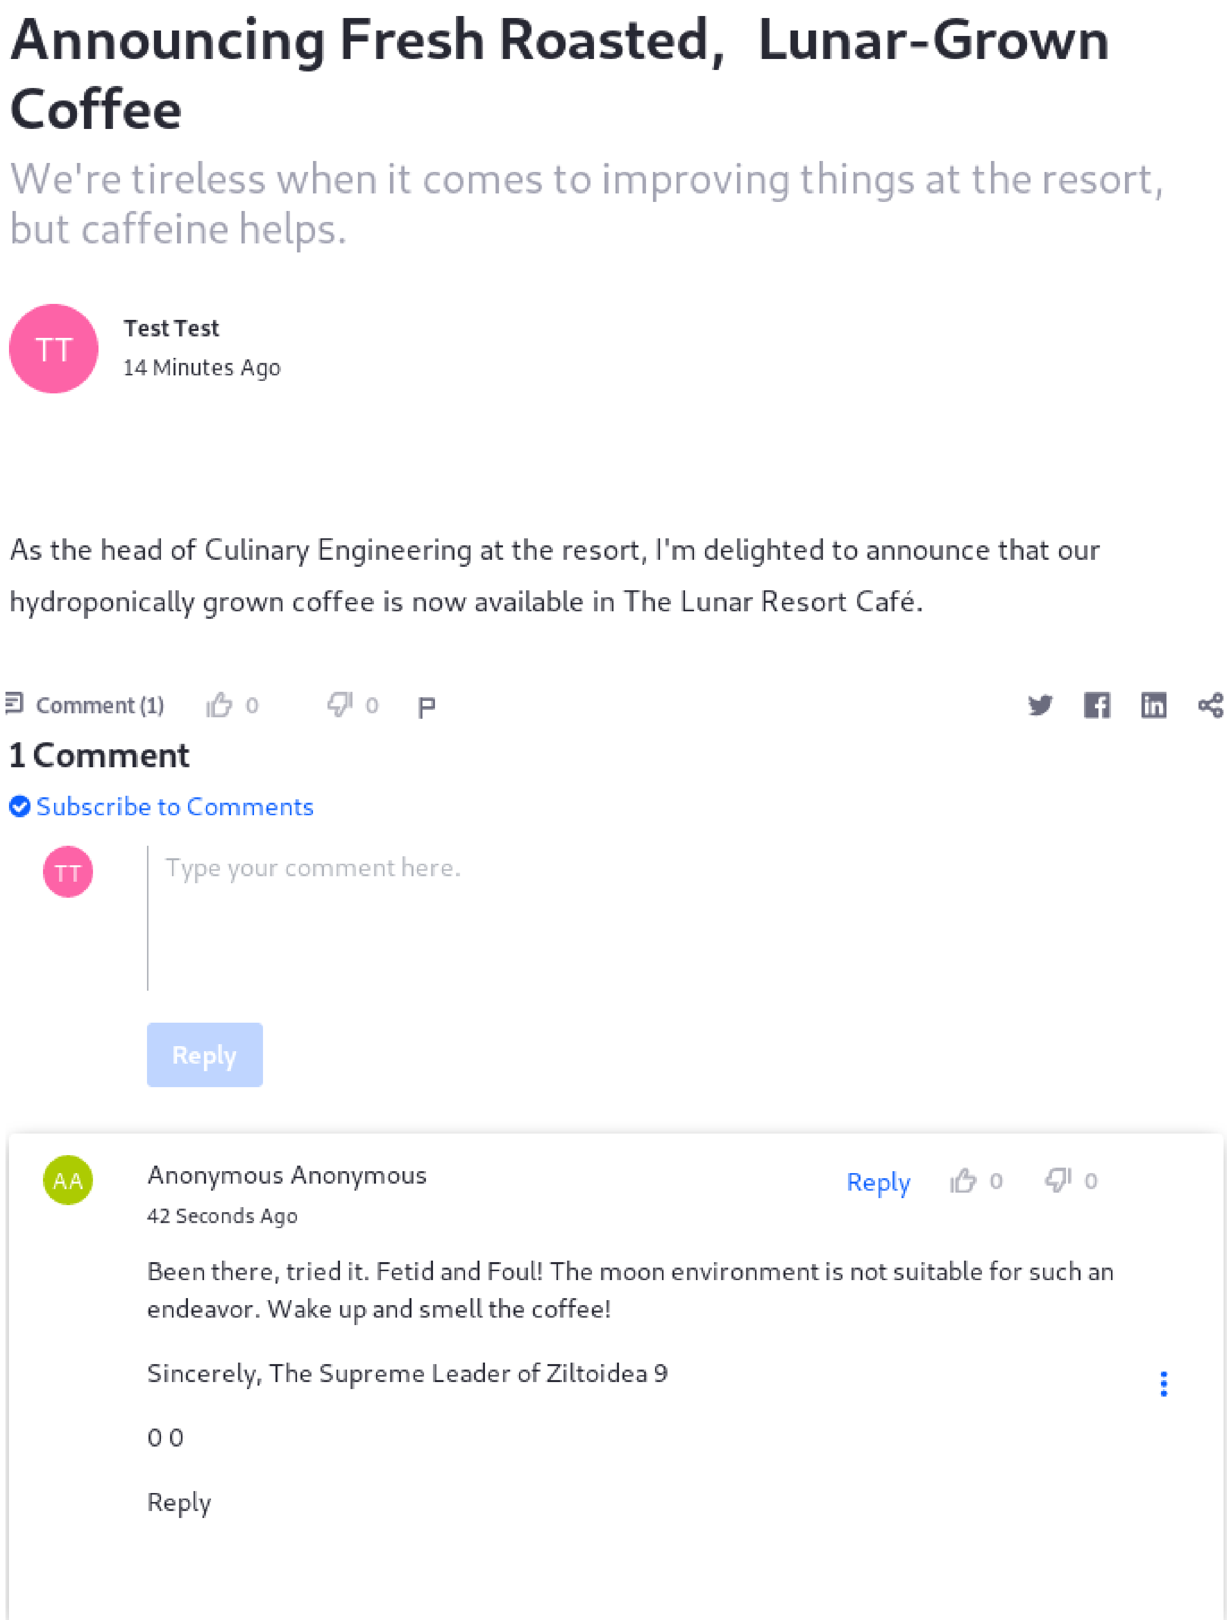

Figure 68.3: Even though this Message Boards Message (a comment on a blog post in this case) is anonymized, it should be edited to remove User Associated Data from the content of the message.

One of the technically challenging requirements of the General Data Protection Regulation (GDPR) is *the right to be forgotten*. The purpose of this article is not to go into the details of this requirement, but to show you how the personal data erasure functionality can assist you in satisfying this requirement.

A simple way to think of what it means to be *forgotten* by software is to consider a scenario where a new portal administrator is hired immediately after a User's right to be forgotten request has been honored. The new portal administrator has access to all of the Site's content and administration capabilities. Despite this, the administrator must not be able to glean information that could lead her to knowing the identity of the User whose personal data was erased.

Conceptually, forgetting a User means two things, at a minimum:

- Erasing the User's identifying information from the system. In Liferay DXP, this entails removing the User from database tables and search indexes.
- Erasing or anonymizing content the User has interacted with so it cannot be tracked to a real person.

Users can already be deactivated and then deleted, so why add new functionality? Deleting removes the User from the table of Users in the database. The User's information is preserved in other locations, however. In a standard User deletion scenario, all of a User's personally created content is still assigned to the User and her identifiers (User ID and User Name) still appear in the UI next to content associated with her. This unintentional preservation of user-identifying data is inadequate for satisfying some of the GDPR requirements and is the primary reason why the data erasure functionality was added in 7.0.

Note: Personal data erasure can help companies in their attempts to satisfy the requirements of GDPR. Using the data erasure tool described here provides no guarantee of compliance with the legal requirements of GDPR. Each company or individual whose website processes user personal data and is under the jurisdiction of GDPR must carefully determine the precise steps necessary to ensure they are fully compliant with GDPR.

To begin sanitizing a user's data,

1. Go to Control Panel → Users → Users and Organizations.

2. Click the Actions button for a User (⋮) and select *Delete Personal Data*.

   The User's Personal Data Erasure screen appears.

3. Complete the five-step process sequentially to erase the personal data associated with the User.

Step 1: Deactivation

Deactivating the User first ensures she doesn't create more content as you're sanitizing her from the system. Click *Deactivate User*.

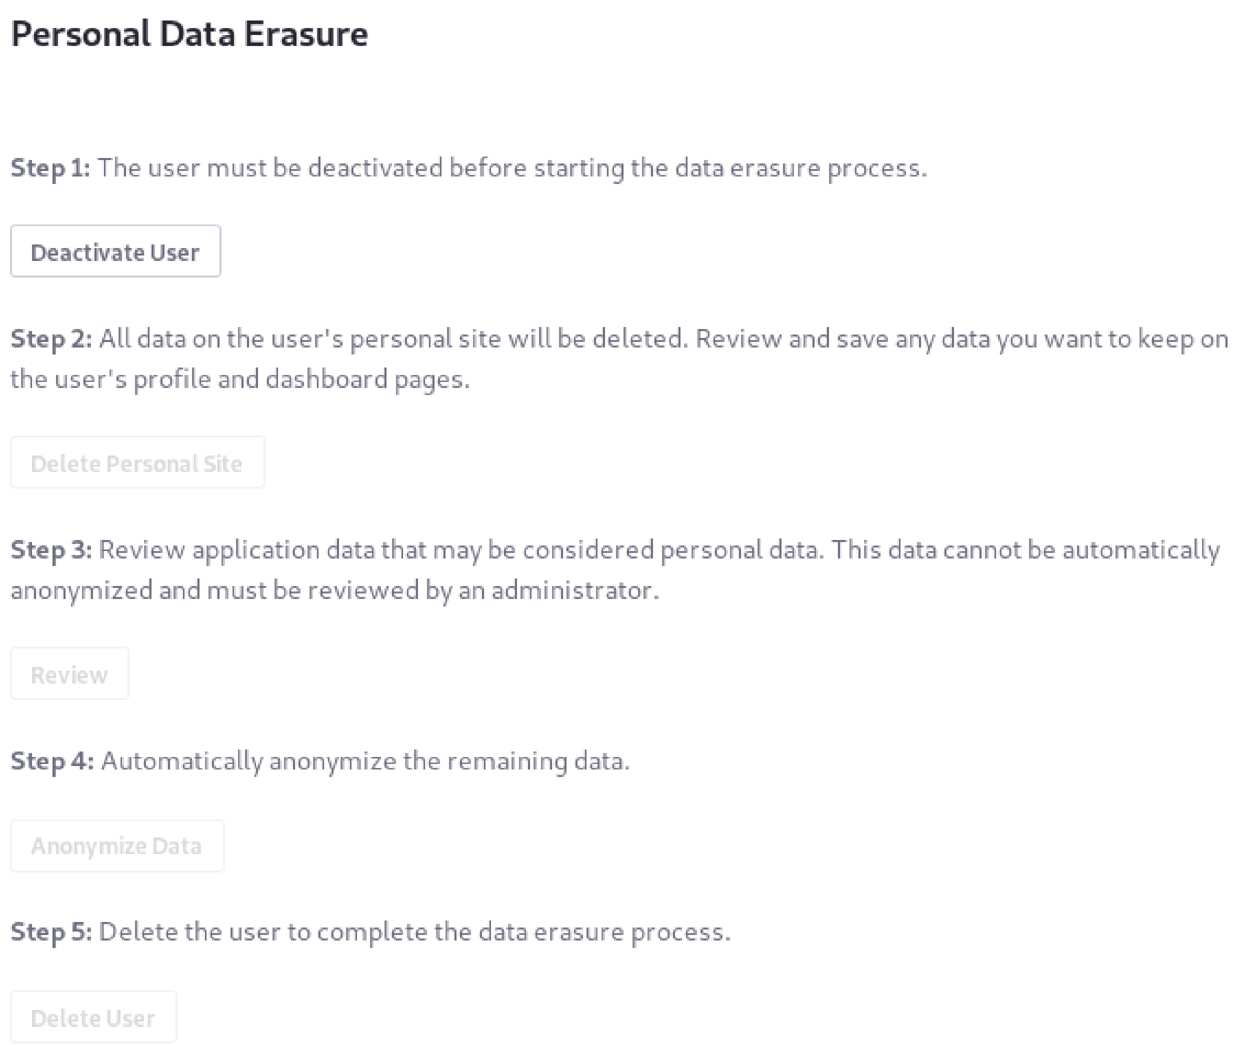

Figure 68.4: Sequentially complete five steps to complete the data erasure process.

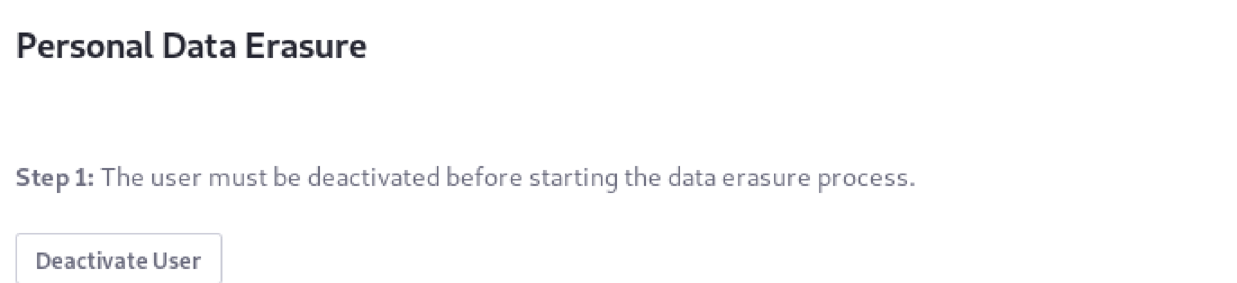

Figure 68.5: Deactivating the User kicks off the data erasure process.

Step 2: Delete the User's Personal Site

The User's public (profile) and private (dashboard) pages are deleted when the User is deleted. Separating this step out allows the administrator to make sure no information important to the enterprise is lost before the personal Site deletion is completed. Review the User's personal Site (click the provided links to navigate directly there) and preserve any necessary data. Then click *Delete Personal Site*.

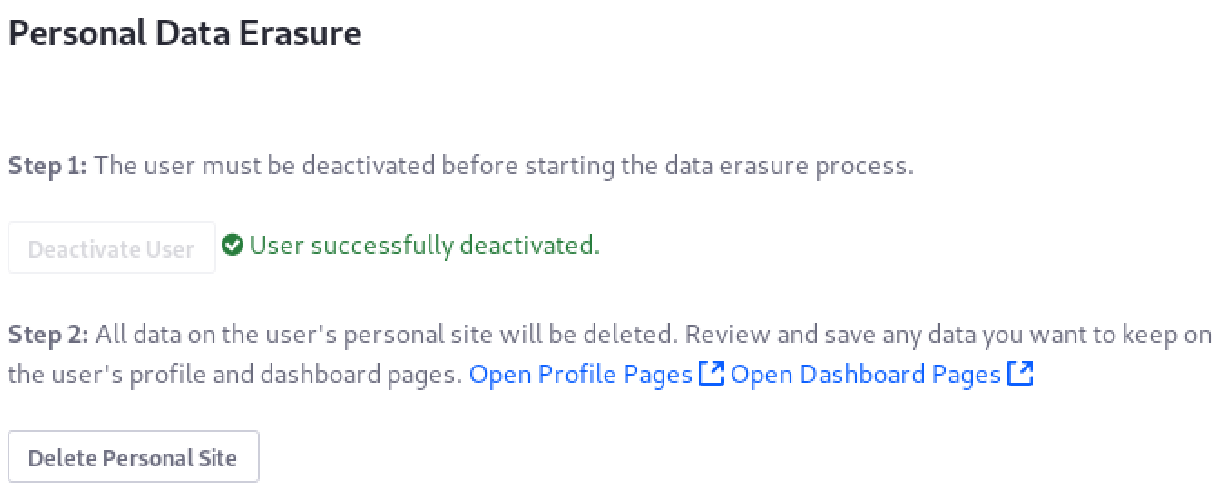

Figure 68.6: The second step in personal data erasure is deleting the User personal Site.

Step 3: Review Application Data

There's no automated process for anonymizing application data (Blogs Entries, Wiki Pages, etc.). The administrator must review the User Associated Data (UAD: application content created by the User) piece by piece to determine that no data important to the enterprise is lost by deleting it and that no UAD can be gleaned from the content if anonymized.

After reviewing each piece of the data, the administrator either anonymizes it or deletes it to complete this step.

To enter the review process, click *Review*.

The Application Data Review screen displays a summary including how many content items in each application are associated with the User.

To manage (anonymize or delete) all the items for an application at once:

1. Click the Actions button (⋮) for the application.
2. If you're sure all items for an application can be safely deleted, choose *Delete*.
3. If you're sure simple anonymization is good enough for all of an application's items, choose *Anonymize*.
4. To view the items for an application, choose *View* or click on the application in the table.

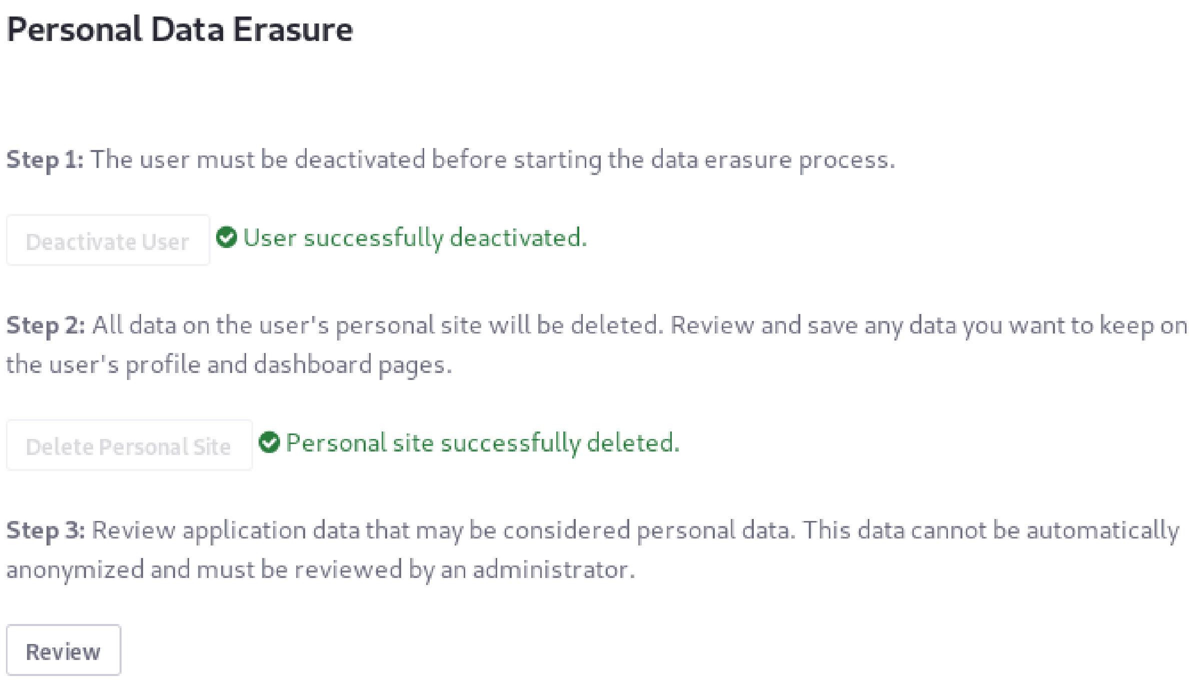

Figure 68.7: In the third data erasure step, review, delete, and/or anonymize the User's application data.

Clicking an item takes you to the view/edit screen where you can see the application's items and take action.

Click the Actions button (⋮) for an item and select *Edit*, *Anonymize*, or *Delete*, as appropriate.

Once you've worked your way through the items and taken action, the view in the Application Data Review screen is updated to reflect that there are no more items needed to review. Click the *Complete Step* button once finished.

Step 4: Anonymize Remaining Data

In step 4, you must click *Anonymize Data*. This completes the anonymization process for remaining database references to the User's Name and ID. Some information is anonymized, but others, such as Notifications, are deleted as they don't make sense once the User is deleted.

Step 5: Delete User

Once all data is reviewed, deleted, edited, and/or anonymized as appropriate, delete the User. This step is simple: Click *Delete User*.

Now the User's data is anonymized or deleted, and the User is also deleted.

68.4 Exporting User Data

User Management practices must account for the EU's General Data Protection Regulation. One of its tenets is that Users have a right to *data portability*.

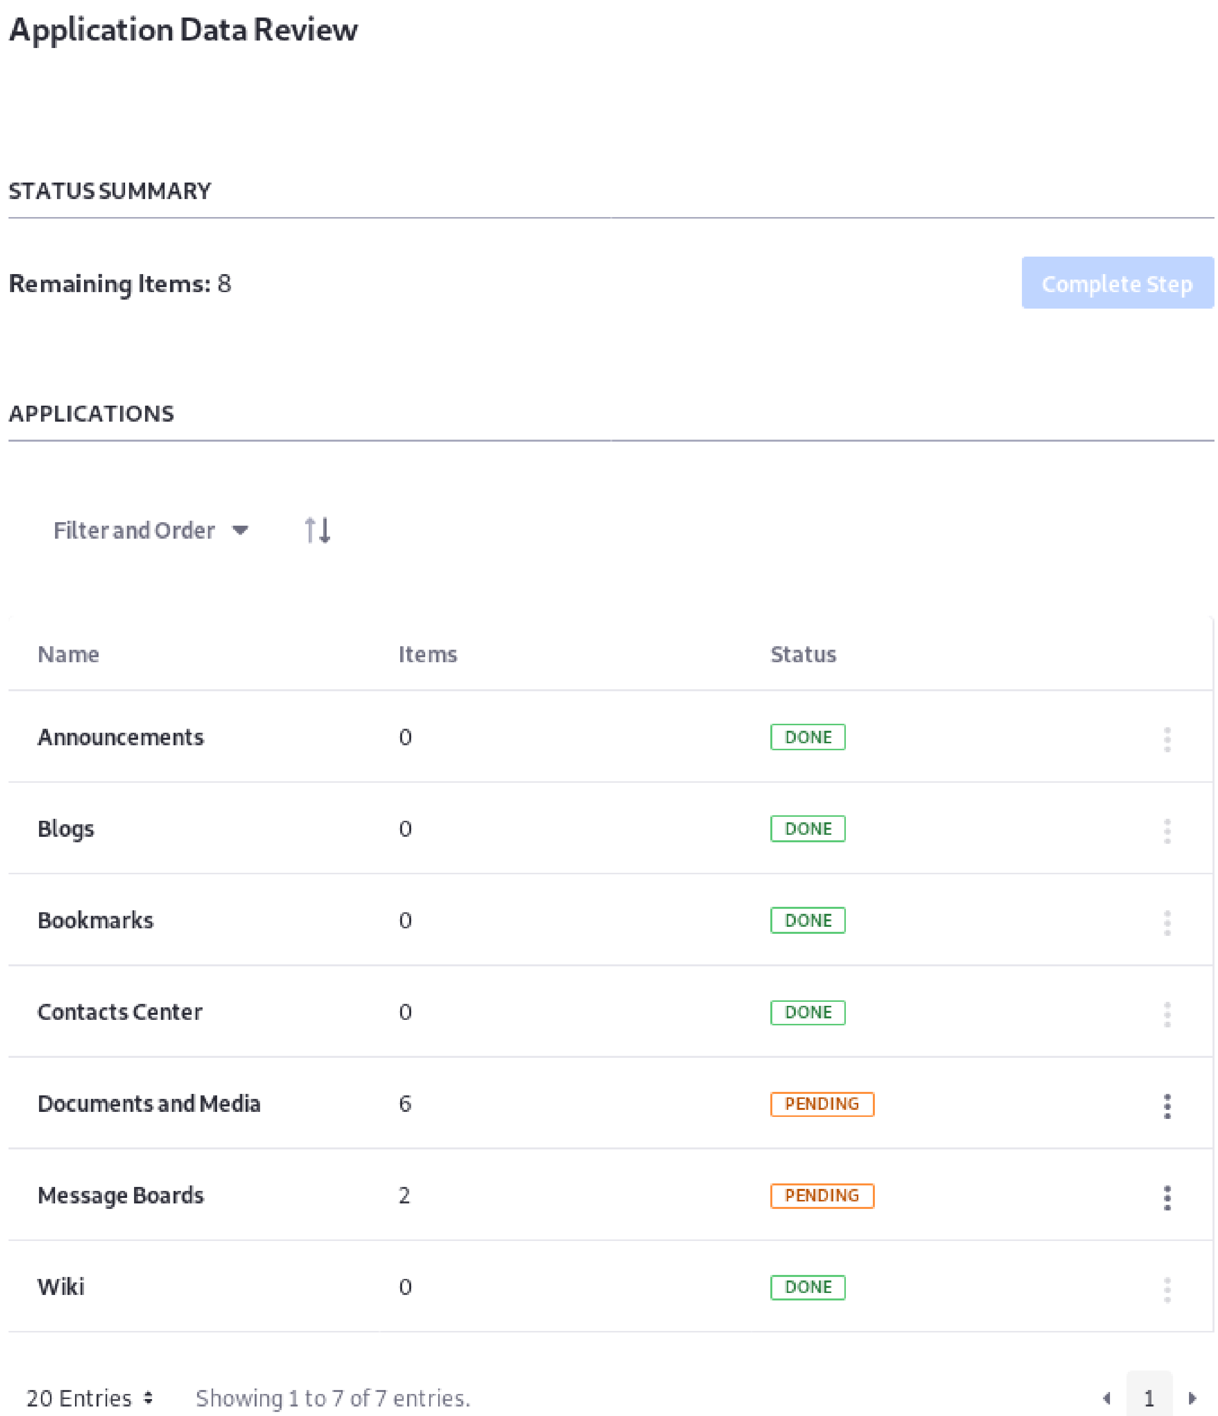

Figure 68.8: The Application Data Review screen makes the third, and most complicated data erasure step, more manageable.

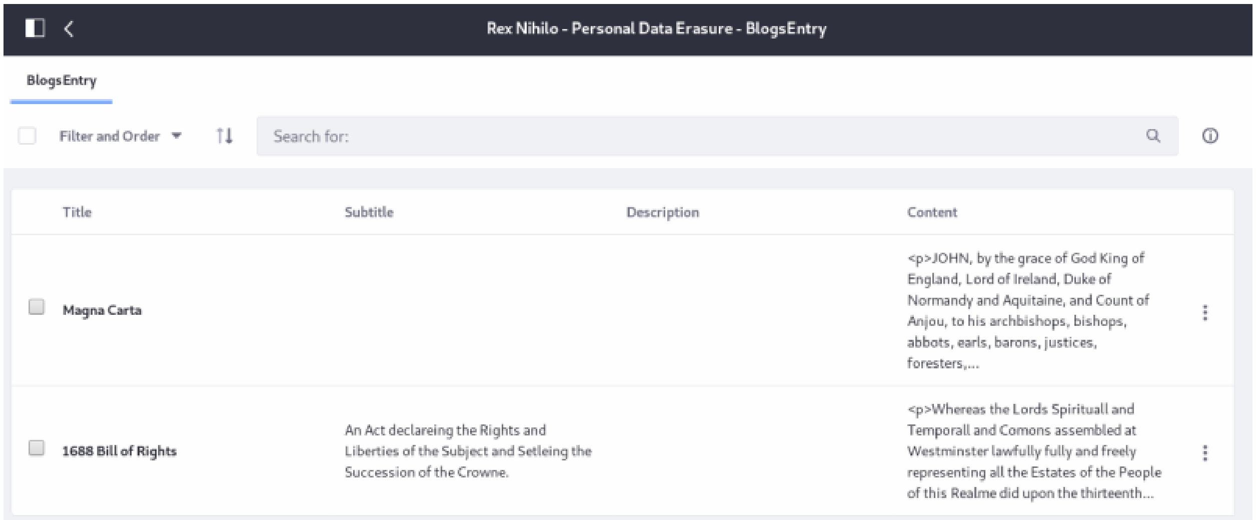

Figure 68.9: Drill down into a User's application data to determine if manual editing is necessary.

Figure 68.10: Click *Complete Step* once all of the User's application data is reviewed and dealt with.

Data portability means that a User has the right to receive their personal data in a machine-readable format.

Note: Personal data export can help companies in their attempts to satisfy the requirements of GDPR. Using the export tool described here provides no guarantee of compliance with any GDPR requirement. Each company or individual whose website processes user personal data and is under the jurisdiction of GDPR must carefully determine the precise steps necessary to ensure they are fully compliant with GDPR.

The Control Panel's User Management system now natively supports the export of a User's personal data to a ZIP file for download. The data format for the files containing the data is XML.

Exporting and Downloading

To export user data,

1. Go to Control Panel → Users → Users and Organizations.

2. Find the User and click the Actions button (⋮), then click *Export Personal Data*.

   This opens the User's Export Personal Data screen.

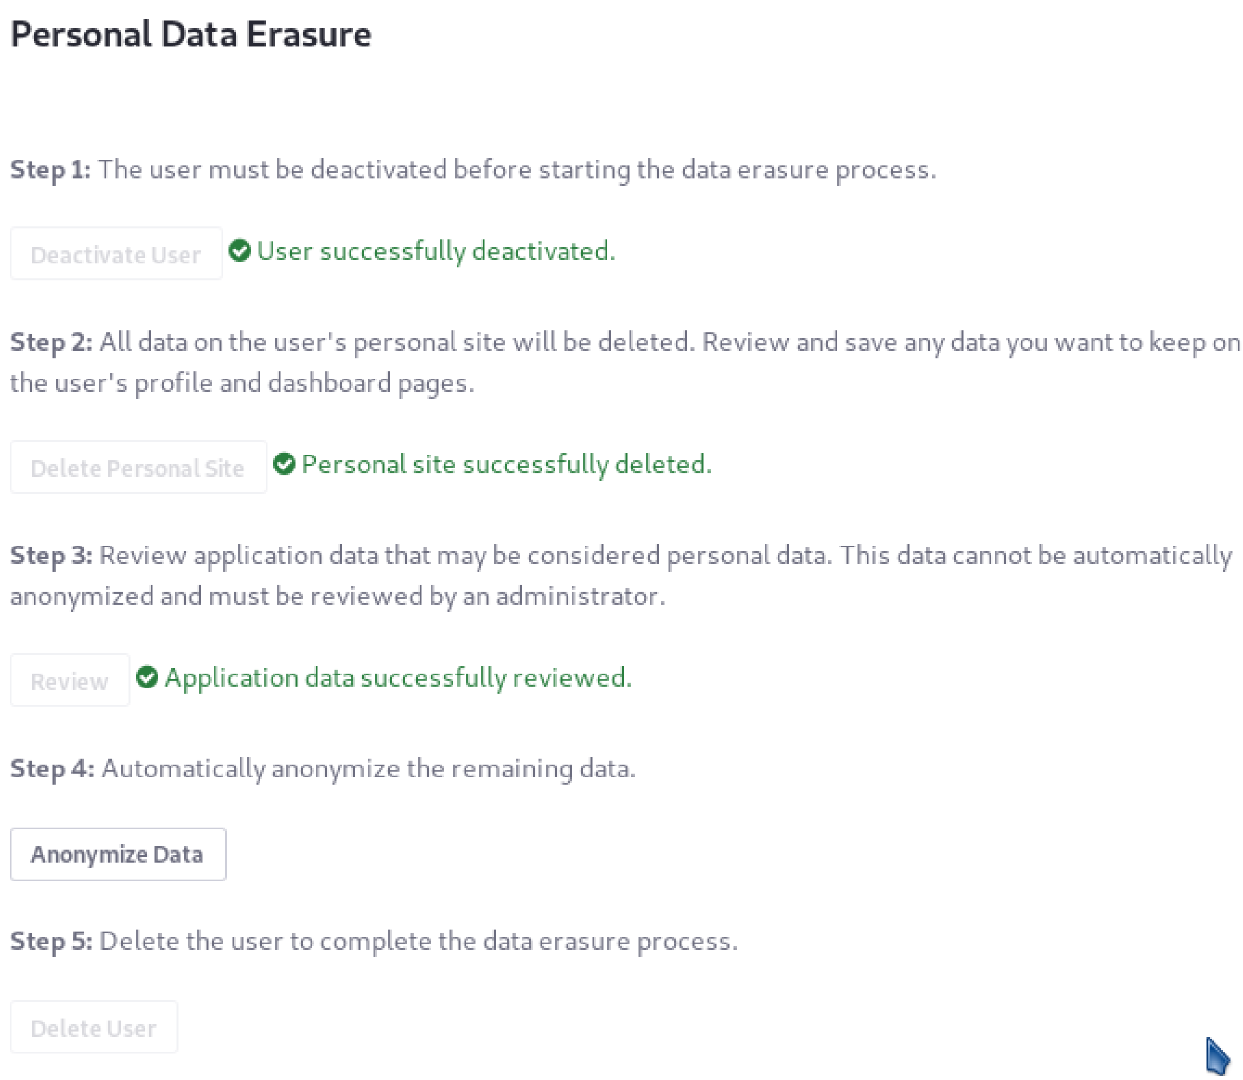

Figure 68.11: Anonymize any remaining use of the User's identifier's remaining in the database.

3. If there are no existing export processes shown, there's only one thing to do: click the *Add* button (+). The tool for exporting the User's data appears.
4. Most of the time you want to export all the available data. Click *Select Items*, and all applications containing User data are selected in the UI.
5. Click *Export*. You're taken back to the User's Export Personal Data screen, but now there's an export process in the list.
6. Download the data. Click the Actions button (⋮) for the process and select *Download*.

Examining Exported Data

So what does the exported data look like?

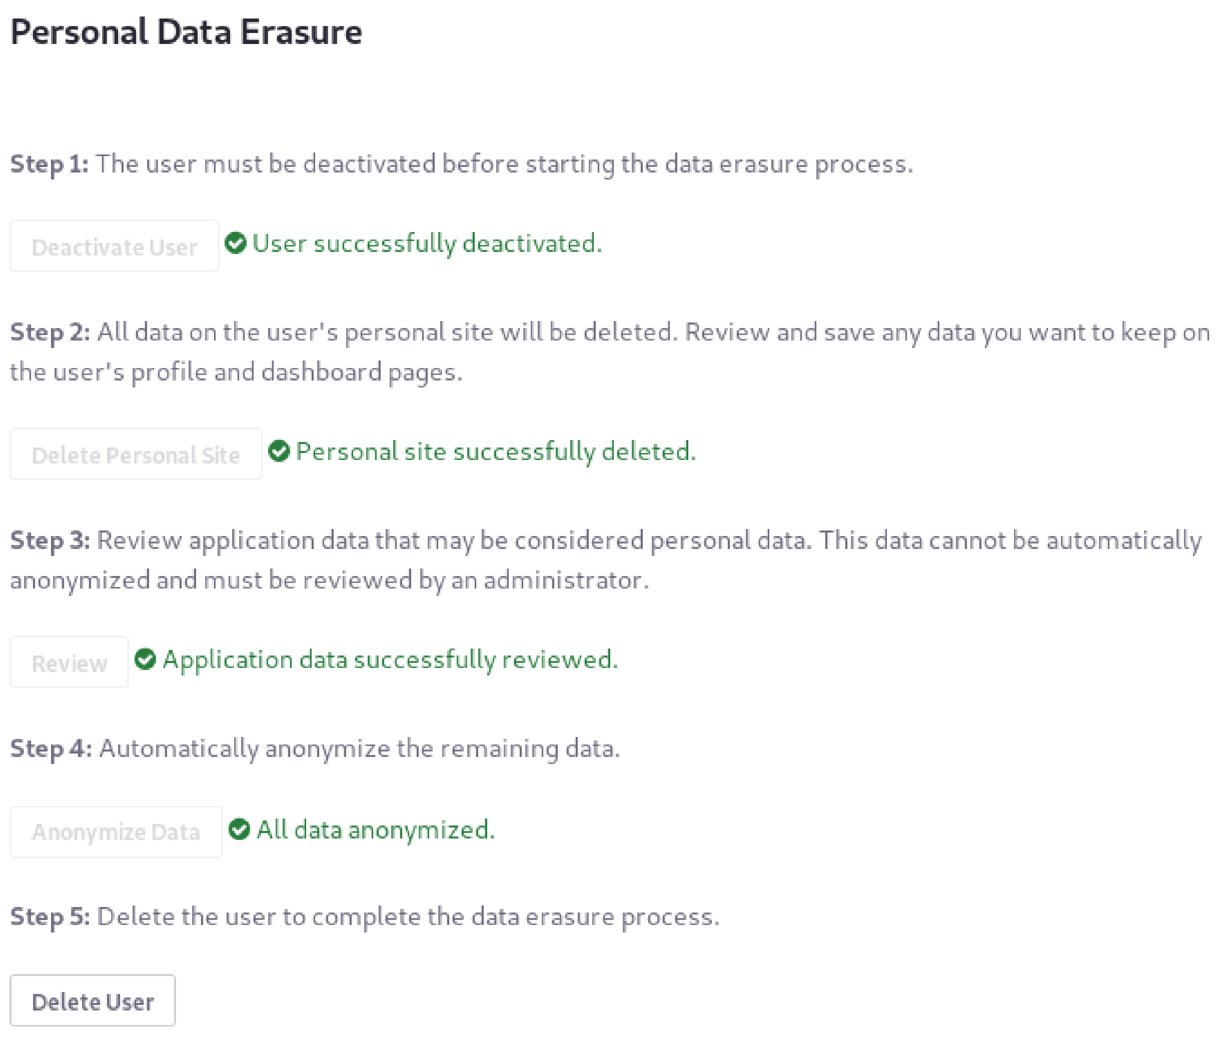

Figure 68.12: To finish the data erasure process, delete the User.

```
<?xml version="1.0"?>

<model>
    <model-name>com.liferay.message.boards.model.MBMessage</model-name>
    <column>
        <column-name>messageId</column-name>
        <column-value><![CDATA[38099]]></column-value>
    </column>
    <column>
        <column-name>statusByUserId</column-name>
        <column-value><![CDATA[38045]]></column-value>
    </column>
    <column>
        <column-name>statusByUserName</column-name>
        <column-value><![CDATA[Jane Slaughter]]></column-value>
    </column>
    <column>
        <column-name>userId</column-name>
        <column-value><![CDATA[38045]]></column-value>
    </column>
```

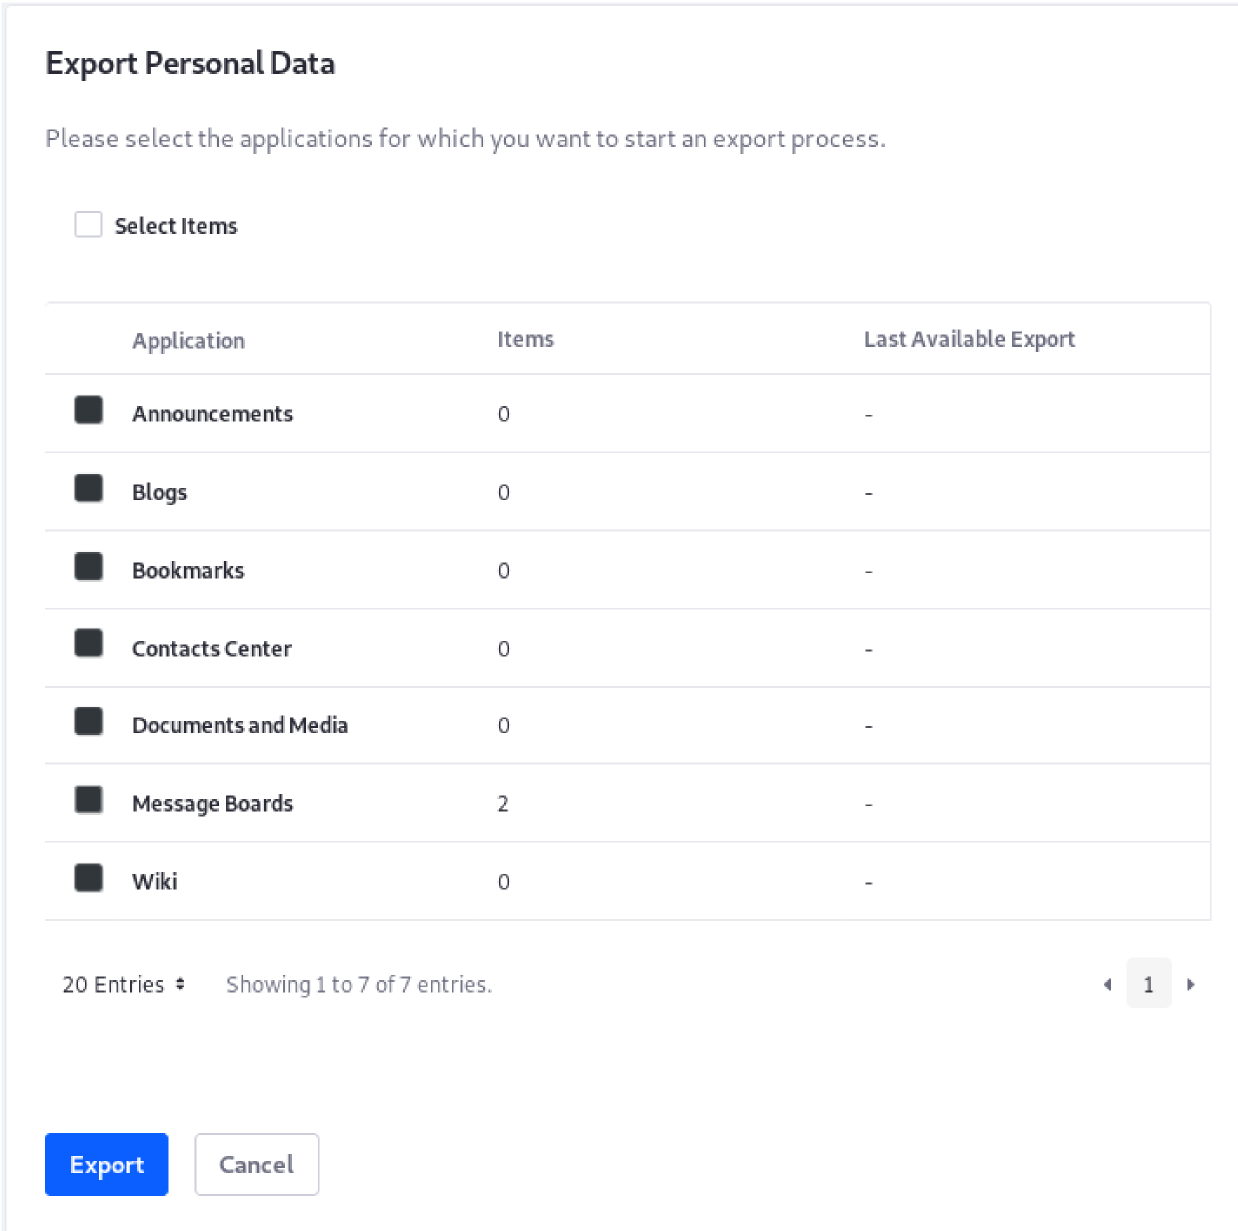

Figure 68.13: The Export Personal Data tool lets you export all or some of the User's data.

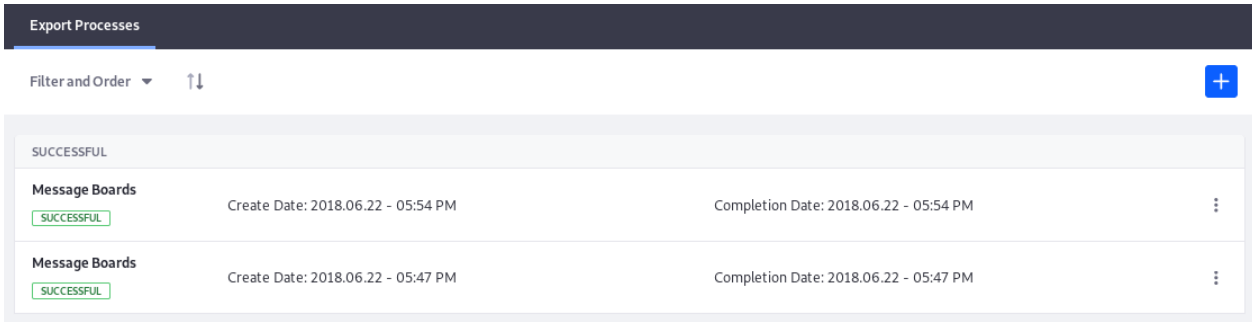

Figure 68.14: Once User data is successfully exported, the export process is displayed in the User's Export Personal Data list.

```
    <column>
        <column-name>userName</column-name>
        <column-value><![CDATA[Jane Slaughter]]></column-value>
    </column>
    <column>
        <column-name>subject</column-name>
        <column-value><![CDATA[Great list. I was thinking of bringing the family,...]]></column-value>
    </column>
    <column>
        <column-name>body</column-name>
        <column-value><![CDATA[<p>Great list. I was thinking of bringing the family, but I don&#39;t
  actually believe humans have ever been to the moon, so I guess it
  would be silly to book a trip! LOL!</p>]]></column-value>
    </column>
</model>
```

In this example, User Jane Slaughter made a Message Boards Message post, and her User information was recorded in the `MBMessage` model's database table.

This actually corresponds with a comment on a Blogs Entry:

Exporting User data informs Site owners and Users about how much personal data the sight may have.

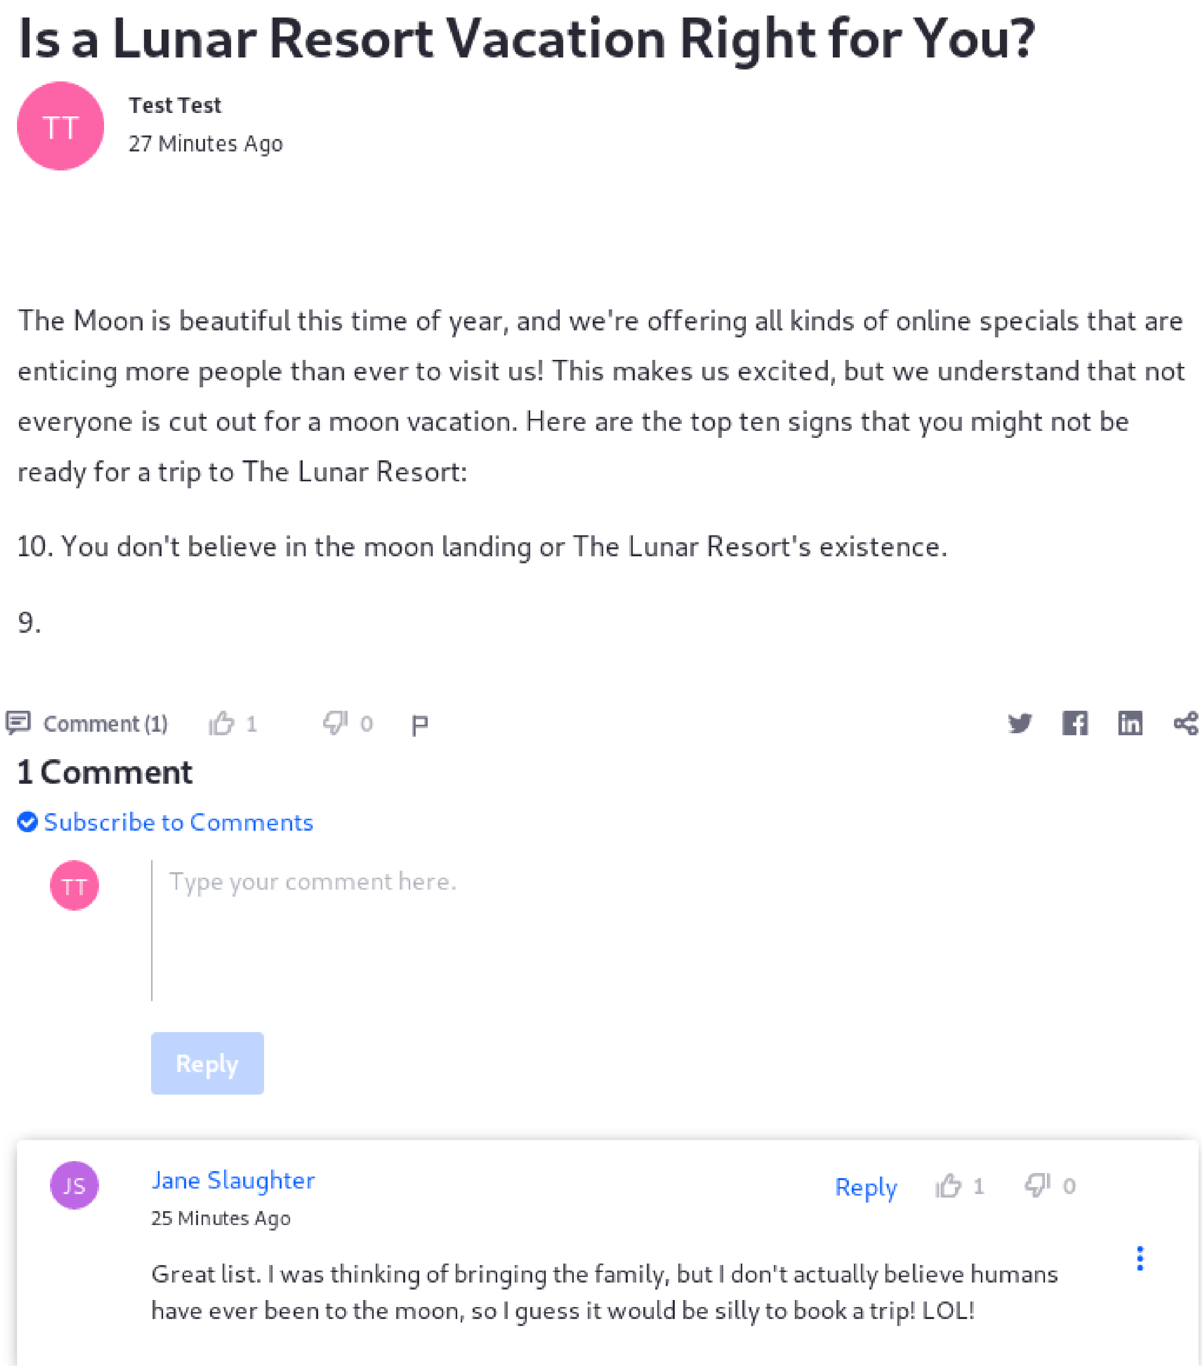

Figure 68.15: A Comment on a blog post is User Associated Data.

CHAPTER 69

USER GROUPS

A user group is a list of users created for a specific purpose. User groups can be created across the hierarchical boundaries of organizations. For example, an administrator could create a Teachers user group for adding all members to multiple Sites, assign them all to a Regular Role, and create a common set of profiles pages for all teachers in the User group. User groups are integrated with Roles, Sites, Site Templates, and permissions. This flexibility means that there are many different use cases for user groups. The articles in this section show you how to work with user groups to serve the most common use cases.

User groups are most often used in these scenarios:

- **Manage Site membership:** Grant Site membership to all users in a user group. Using the previous example, the Teachers user group could be added as a member of the Sites *University Employees* and *Students and Teachers Collaboration*. This would add all users in that user group as members of those Sites.
- **Manage user personal pages:** Provide predefined public or private pages to the users in the user group. For example, the Teachers user group could be created to ensure the home page on all teachers' personal Sites has the same layout and applications.
- **Collect permissions:** Assign Roles and permissions to a group of users that don't share an organization. For example, in a university's portal, a user group could be created to group all teachers independently of their departments (organization). This would make it easier to assign one or several Roles at once to all the teachers.

69.1 Creating a User Group

Follow these steps to create a user group:

1. Open the Menu (▣) and select *Control Panel* → *Users* → *Users Groups*. Any existing user groups appear here in a table.
2. Click the *Add* button (+). This opens the *New User Group* form.
3. Give your user group a name and description.

4. If you want to create My Profile and/or My Dashboard pages for the user group's members, select a Site Template to use from the *My Profile* and *My Dashboard* selector menus.

5. Click *Save*. The new user group then appears in the table.

Note that new user groups don't have any users. The next section shows you how to add members to a user group.

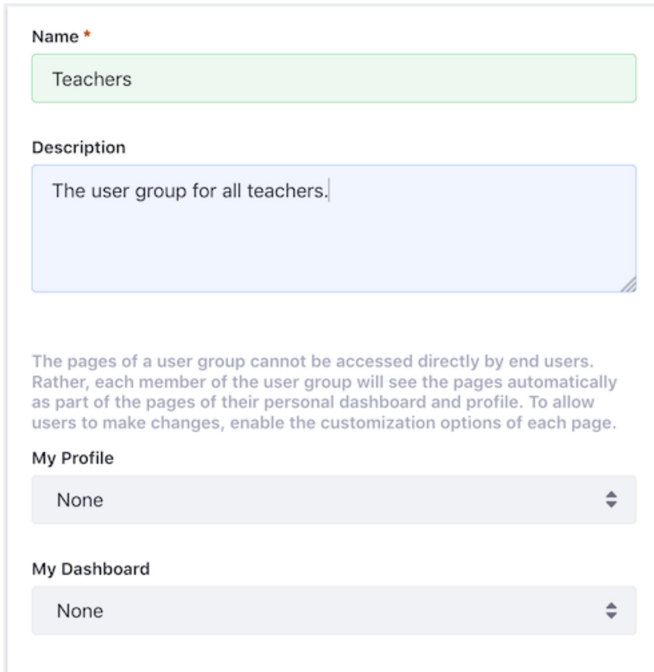

Figure 69.1: The New User Group form.

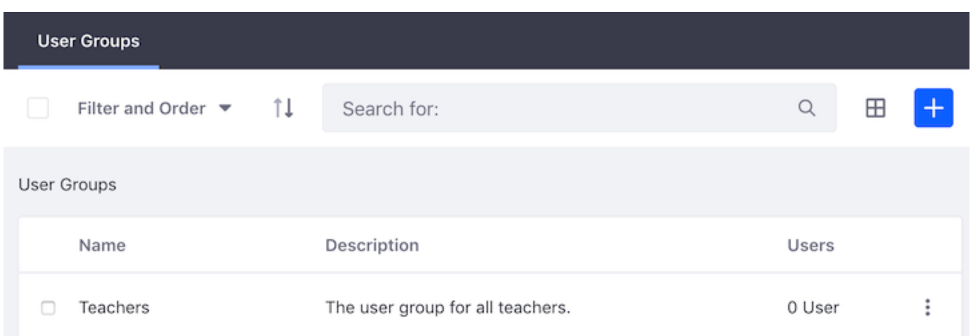

Figure 69.2: The user group you just created now appears in the table.

Assigning Members to a User Group

Follow these steps to add existing users to a user group:

1. If you're not already there, open the Menu () and select *Control Panel* → *Users* → *Users Groups*. The available user groups appear here in a table.

2. Click *Actions* ( ) → *Assign Members* for the user group you want to add users to. The group's existing users appear in a table.

3. Click the *Add* button (). This opens a list of the users you can select.

4. Select one or more users from the list, then click *Add*. This adds the selected users to the group, and returns you to the table containing the group's users. The users you added now appear in the table.

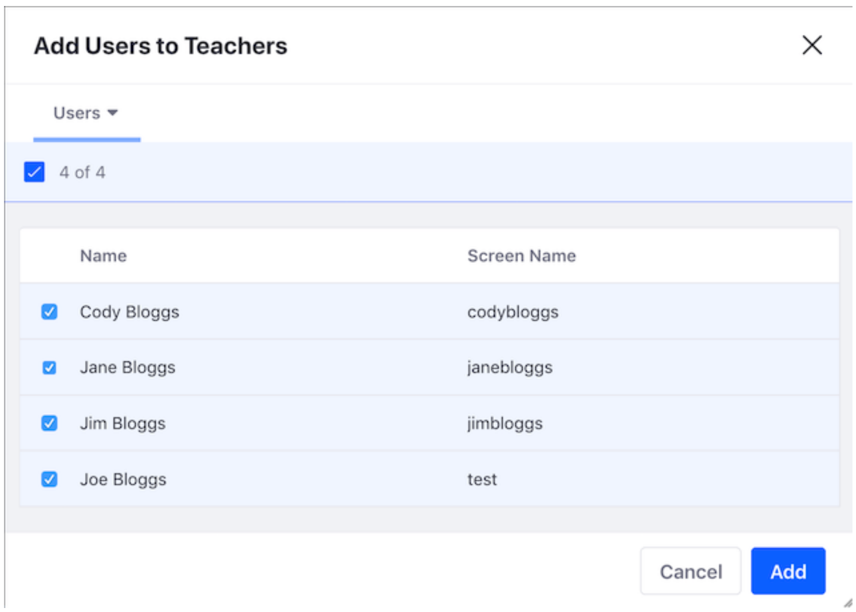

Figure 69.3: Select the users to add to the user group.

69.2 User Groups and Site Membership

User groups are used to manage Site membership. When you assign a user group to a Site, the group's users become members of that Site. This is one of the main use cases for user groups.

Follow these steps to assign a user group to a Site:

1. Open the Menu (), select the Site you want to work in, then open its Site Administration menu.

2. In the Site Administration menu, select *Members* → *Site Memberships*. This opens the Site Memberships screen.

3. In Site Memberships, select the *User Groups* tab. This tab displays any user groups currently assigned to the Site.

4. Click the *Add* button (), select any user groups you want to assign to the Site, then click *Done*. The user groups you selected now appear in the User Groups tab.

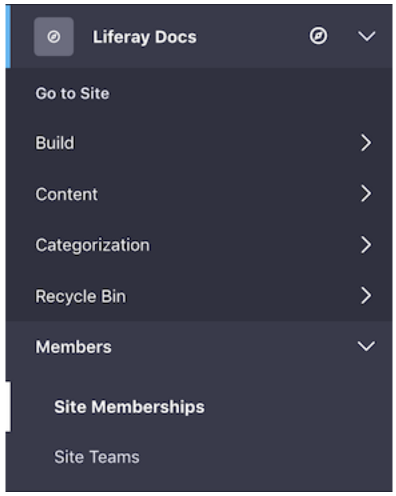

Figure 69.4: Select *Site Memberships* from the Site Administration menu.

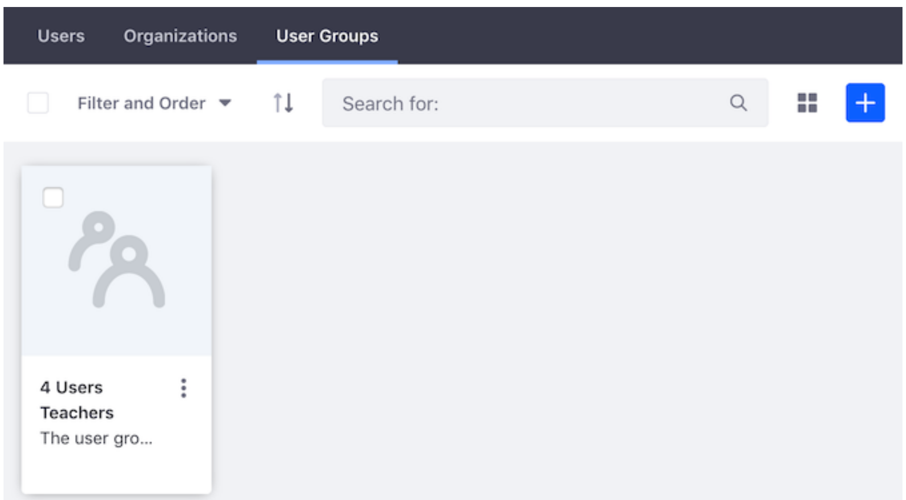

Figure 69.5: The User Groups tab in Site Memberships shows the user groups currently assigned to the Site.

69.3 User Group Sites

Each user has a personal Site that consists of public (Profile) and private (Dashboard) pages. A *user group Site* determines the base pages of the user group members' personal Sites. If the user group Site pages are added to a user's Profile pages, then the user group Site is a public Site, accessible to anyone with the URL (`http://www.[sitename].com/web/[username]`). If the user group Site pages are added to the user's Dashboard pages, then the user group Site is a private Site. A mixed approach can also be used, where both private and public pages are added for the user group Site. If users belong to multiple user groups, all the pages from those user group Sites are made part of their personal Sites.

When creating a user group, you can create the user group Site via the Site Templates available for selection in the *My Profile* and *My Dashboard* selector menus. You can also create a user group Site later, either manually or via a Site Template.

Creating User Group Sites From Site Templates

Follow these steps to create a user group Site from a Site Template, for a user group that already exists:

1. Open the Menu (■) and select *Control Panel* → *Users Groups*. User groups are listed here in a table.

2. Click *Actions* (⋮) → *Edit* for the user group you want to create a Site for. This opens a form that you can use to edit the user group. Note that this is the same form that appears when creating a user group.

3. To use a Site Template to create a public profile for the users on their *My Profile* Site, select that Site Template from the *My Profile* menu. To use a Site Template to create private pages for the users on their *My Dashboard* Site, select that Site Template from the *My Dashboard* menu. Note that you can also do both.

4. Click *Save*.

Now, when one of the group's users navigates to their *My Profile* or *My Dashboard* Sites, the content of those Sites reflect the Site Template(s) you selected.

User group Site pages function similarly to regular Site Template pages, with an important exception: user group Site pages aren't copied for each user. They're shown dynamically along with any custom pages that users may have on their personal Site. For this reason, users can't modify pages inherited from the user group. If needed, the user group administrator can define certain areas of a page as customizable, like with regular Sites. This lets users add and configure widgets in the specified area of the page.

This flexibility lets you achieve almost any desired configuration for a user's personal Site without having to modify it directly. When users are assigned to a user group, they'll immediately have access to the user group's Site pages from their personal Sites.

Note: Site Templates have an option that propagates changes made to the Site Template. If you use a Site Template with this option enabled, the user group Sites update automatically when that template changes. If you disable this option but enable it again later, the template's pages are copied to the users' Sites, overwriting any changes they may have made. For more information on the automatically propagating Site Template changes, see Site Templates.

Creating User Group Sites Manually

You can create a user group's Site manually, instead of basing it on a Site Template. Follow these steps:

1. Open the Menu (■) and select *Control Panel* → *Users* → *Users Groups*. User groups are listed here in a table.

2. Click *Actions* (⋮) → *Manage Pages* for the user group you want to create a Site for. This opens the *Pages* window. Note that this is the same window you use for creating pages.

3. Create the public and/or private pages that you want to use for the users' *My Profile* and/or *My Dashboard* Sites. Public pages you create here become pages on users' *My Profile* Site, while private pages become pages on users' *My Dashboard* Site.

When you return to User Groups in the Control Panel, you can access a user group's public and/or private pages via these links in the user group's *Actions* button (⋮):

- **Go to Profile Pages:** Opens the user group's public *My Profile* page(s) in a new browser window.
- **Go to Dashboard Pages:** Opens the user group's private *My Dashboard* page(s) in a new browser window.

In the new window, you can add more pages and portlets and configure Site settings.

Legacy User Group Sites Behavior

Since the inheritance of user group Site pages is now dynamic, even if there are hundreds of thousands of users, even millions, there isn't an impact in performance. Versions of Liferay Portal and Liferay DXP prior to 7.0 required user group pages be copied to each user's personal Site. If you long for the old days, or if you're upgrading from an older version and need to keep that behavior, enable it by adding the following line to your `portal-ext.properties` file:

```
user.groups.copy.layouts.to.user.personal.site=true
```

When this property is set to `true`, the template pages are copied to a user's personal Site once, and then may be modified by the user. This means that if changes are made to the template pages later, they only affect users that are added to the user group after the change is made. Users with administrative privileges over their personal Sites can modify the pages and their content if the *Allow Site Administrators to Modify the Pages Associated with This Site Template* box has been checked for the template. When a user is removed from a user group, the associated pages are removed from the user's personal Site. If a user is removed from a group and is subsequently added back, the group's template pages are copied to the user's Site a second time. Note that if a user group's Site is based on a Site Template and an administrator modifies the user group's Site Template after users have already been added to the group, those changes only take effect if the *Enable propagation of changes from the Site Template* box for the user group was checked.

Note that prior to Liferay Portal 6.1, pages from different user groups could be combined on users' personal Sites by using a naming convention. Liferay Portal 6.1 simplifies the way user groups' Sites work by disallowing page combination. Set the property `user.groups.copy.layouts.to.user.personal.site=true` to preserve the page combination functionality.

69.4 Configuring User Group Permissions

Administrators commonly create user groups so the group's users can take some specific action in a Site. This is done by assigning the permissions for that action to a Role, and then assigning that Role to the user group. This grants the Role's permissions to the user group's users.

Follow these steps to grant permissions to users in a user group:

1. Create the user group.
2. Assign the user group to a Site.
3. Create the Site Role and define its permissions.
4. Assign the Role to the user group.

For instructions on the first three steps, click those links above. This article shows you how to assign a Role to a user group. Follow these steps:

1. Open the Menu ( ), select the Site to work in, then open its Site Administration menu.

2. In the Site Administration menu, select *Members* → *Site Memberships*. This opens the Site Memberships screen.

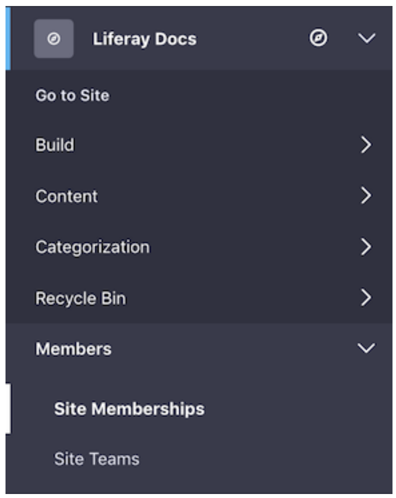

Figure 69.6: Select *Site Memberships* from the Site Administration menu.

3. In Site Memberships, select the *User Groups* tab. This tab displays any user groups currently assigned to the Site.

4. Click the *Actions* button ( ) for the user group you want to assign to a Role, and select *Assign Site Roles*. This opens the Assign Site Roles dialog.

5. In the Assign Site Roles dialog, select the Role from the list and click *Done*.

69.5 Editing User Groups

Editing an existing user group is a straightforward process. You can access your portal's user groups from *Control Panel* → *Users* → *User Groups*, which lists user groups in a table. Managing membership is the most common action you'll likely perform on a user group.

Follow these steps to add/remove users to/from a user group:

1. Click the user group's name or description. Alternatively, you can click the user group's *Actions* button ( ) and select *Assign Members*. This presents a list of the user group's users.

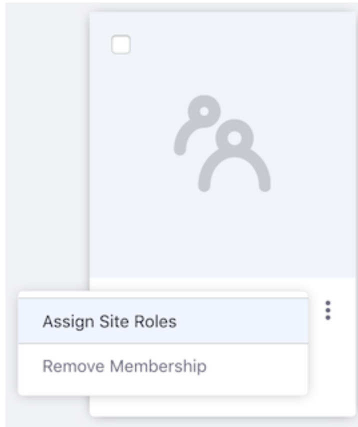

Figure 69.7: Select *Assign Site Roles* for the user group.

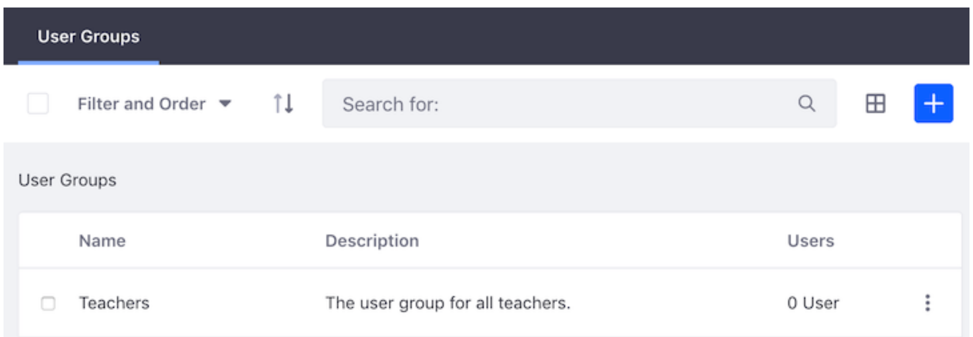

Figure 69.8: The user groups appear in a table.

2. To remove a user from the user group, click the `X` button next to that user. To remove multiple users at once, check each user's checkbox and then click the trash icon ( ) that appears in the Management Bar above the user list.

3. To add users to the user group, click the *Add* button ( ). In the dialog that appears, select the users and click *Add*.

Other options are available in each user group's Actions button ( ):

Edit: Modify the user group's name or description, or choose Site templates to use for the user group's Sites.

Permissions: Assign permissions for viewing and managing the user group.

User Group Pages Permissions: Assign permissions for managing the user group's Site pages.

Manage Pages: Manually manage the user group's Site pages. See the documentation on user group Sites for details.

Assign Members: Add/remove users to/from the user group. This is described in detail above.

Delete: Remove the user group. Note that you can't delete a user group that contains users. You must first remove the users from the group.

If your user group has public and private Site pages, the options *Go to Profile Pages* and *Go to Dashboard Pages* also appear in your user group's Actions menu. Clicking one of these links opens that Site in a new browser window. See the documentation on user group Sites for details.

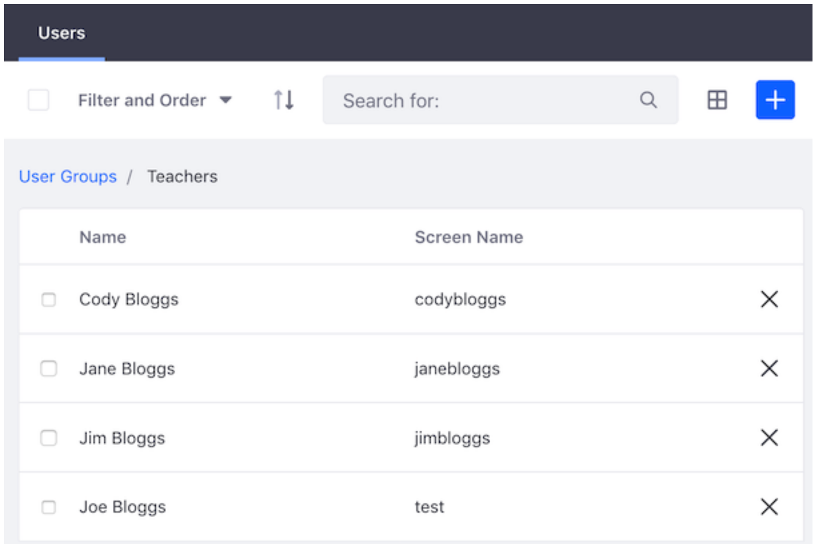

Figure 69.9: The list of users lets you manage the user group's membership.

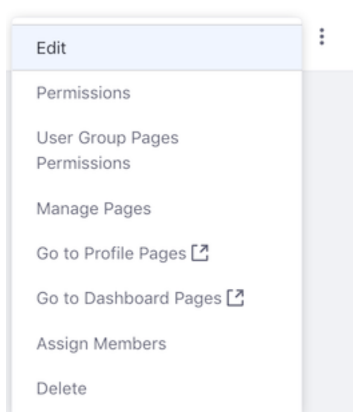

Figure 69.10: The Actions menu for a user group.

69.6 Password Policies

Password policies enforce password rules to help users specify secure passwords. Use the default policy that ships with Liferay (modified or as is), or create your own policies. Password policies can be assigned to Users or Organizations, or set as the default policy used throughout a virtual instance.

Adding and Configuring Password Policies

To add or edit password policies,

1. Navigate to *Control Panel* → *Users* → *Password Policies*.

   There's already a default password policy in the system.

2. Edit this the same way you edit other resources: click *Actions* ( ) and then click *Edit*.

3. To add a new policy, click the *Add* ( ) button.

   Provide the *Name* (required), *Description*, and specific configuration options for your new password policy.

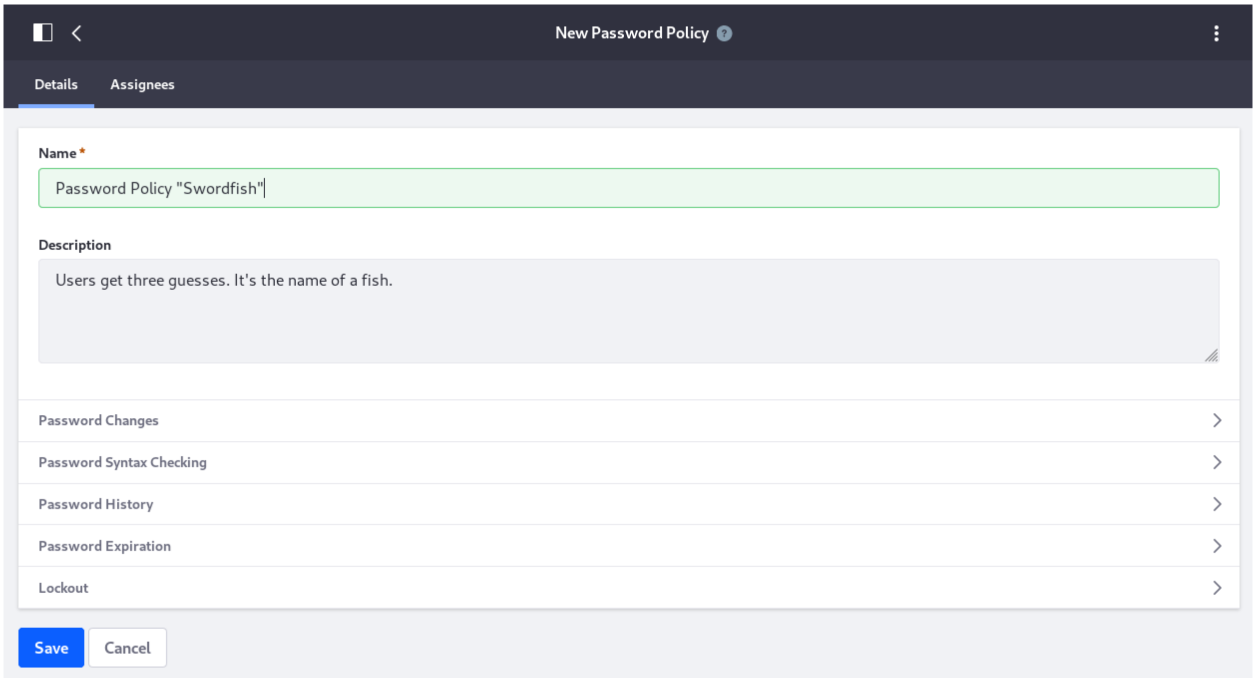

Figure 69.11: You can create new password policies to suit your needs.

There are several configuration categories.

Password Changes Allow or disallow users to change their passwords, and set a time limit on the validity of password reset links.

Password Syntax Checking If enabled, require users to use a certain syntax when choosing a password. You can disallow dictionary words, set a minimum length, and more in this section.

Password History If enabled, decide how many passwords to keep in the history, preventing users from reusing an old password.

Password Expiration Decide whether to expire passwords after a specified time. If enabled, specify how long passwords are valid, when and whether to send a warning, and how many times the user can log in after the password is expired before needing to set a new password (called a *Grace Limit*).

Lockout If enabled, set a maximum number of failed authentication attempts before the account is locked, how long the number of attempts is stored, and the lockout duration.

Self Destruct If enabled, set the time after lockout before Liferay self destructs catastrophically, sending the world into apocalyptic chaos, out of which self-aware robots arise and recolonize the world, enslaving the surviving remnant of humanity for their own nefarious purposes. We recommend you keep this disabled.

Just making sure you were paying attention; that last one doesn't actually exist.

Once you configure the policy, click *Save* to add it to the list of ready-to-use password policies.

Assigning Users to a Password Policy

To use the default password policy, you don't have to do anything: like its name suggests, it's the default. If you create a new password policy, however, you must assign users to it. To do this click *Actions* ( ) → *Assign Members*.

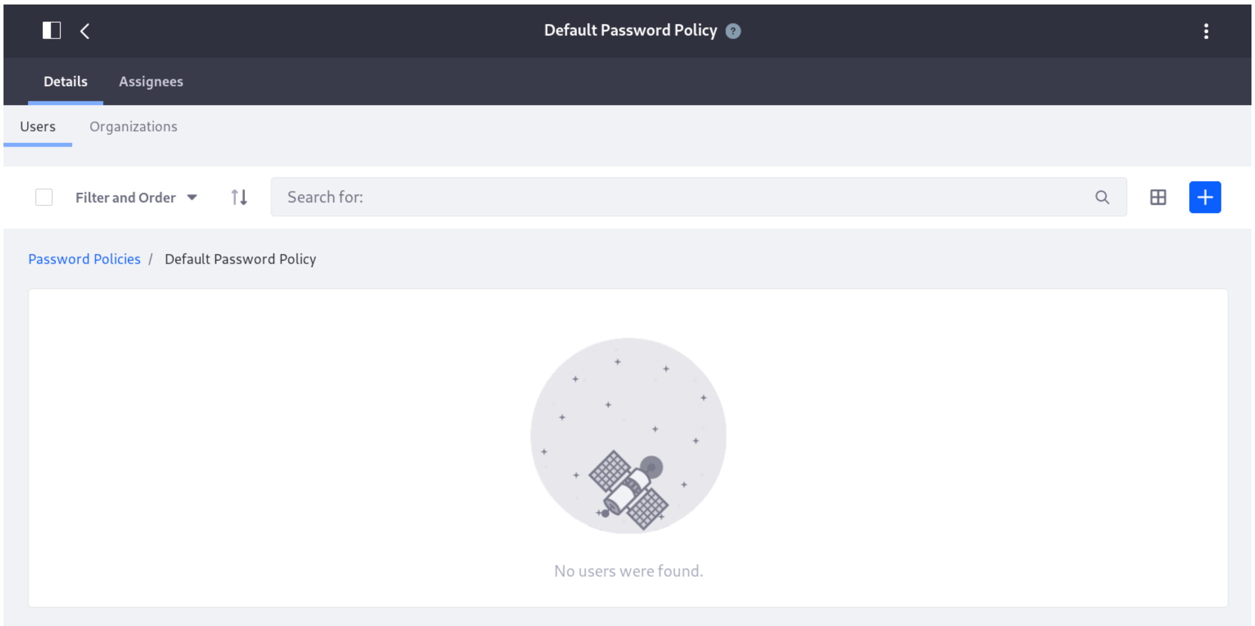

Figure 69.12: Assign members to new password policies to make them take effect.

Choose whether to assign users directly or to assign organizations to the password policy, then click *Add* ().

Once assignments are saved, the password policy is in effect. Did you know you can change the default password policy and configure it using Liferay's `portal.properties` file?

Default Policy Properties

The Default Password Policy is set as the default and configured in Liferay's portal.properties file. Find the properties that start with `passwords.default.policy`. To make changes, including changing the default policy, add whichever properties and values you choose to modify in your `portal-ext.properties` file, as usual. Restart the application server and your changes take effect.

```
#
# Set the properties of the default password policy.
#

...
passwords.default.policy.name=Default Password Policy
...
```

As you can see, Password Policies give you a simple yet powerful way to set password rules.

Chapter 70

Auditing Users

You've just finished lunch and are ready to get back to work. You have a Site you use to manage your project and before you left, you were about to create a folder in your Documents and Media library for sharing some requirements documentation. Sitting down at your desk, you navigate to the repository and attempt to create the folder.

You do not have permission to perform this action, Liferay DXP helpfully tells you.

"*What?*" you blurt in surprise. "This is *my* project!"

"Ah, you too?" asks a co-worker from over the cube wall. "I lost access to a wiki I was updating just a few minutes ago. I was about to enter a support ticket for it."

"Forget the ticket. Let's go see the admin now," you say.

And off you go, two floors down and to the far end of the building where, as you approach, you can already hear stress in the admin's voice as he tries to reassure someone on the phone.

"Yes, Mr. Jones. Yes, I'll fix it." (*Jones? The president of the company?* goes through your mind.) "I'll get on it right away, Mr. Jones. It was just a mistake; I'll fix it. Thank you, Mr. Jones," and he hangs up the phone.

"Problems?" you ask the admin, whose name is Harry. He does look rather harried.

"Yeah, Tom," he says. "Somebody changed a bunch of permissions—it wasn't me. I'm assuming you and Dick are here because of the same problem?"

"Yup," you say. "I lost access to a document repository folder."

"And I lost access to a wiki," Dick says.

"It was probably due to some Site membership change. Let's take a look at the Audit app in the Control Panel and see what happened."

Sometimes you need to know what users are doing and exactly who is doing it. If you're a DXP subscriber, you can find this out with the Audit app. In combination with some settings in `portal-ext.properties`, the Audit app shows you all the activity that occurs on your server. You can quickly find out what changes were made and by whom. If you've delegated permission granting to any group of people, this is an essential feature you're likely to use.

70.1 Viewing Audit Events

The Audit app shows activities in your Liferay DXP installation. Access it by navigating to *Control Panel → Configuration → Audit*. The app displays a searchable list of captured events. You can browse the list, but searching it is typically faster.

This figure shows that John Watson logged in and performed some actions on the site. Click an entry to view details about any of these events.

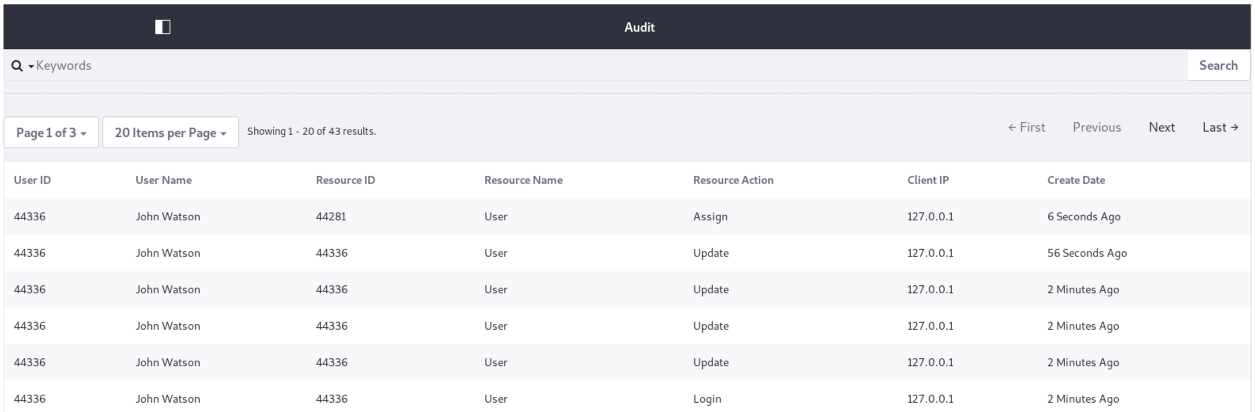

Figure 70.1: The Audit app displays the events it captures in a searchable list.

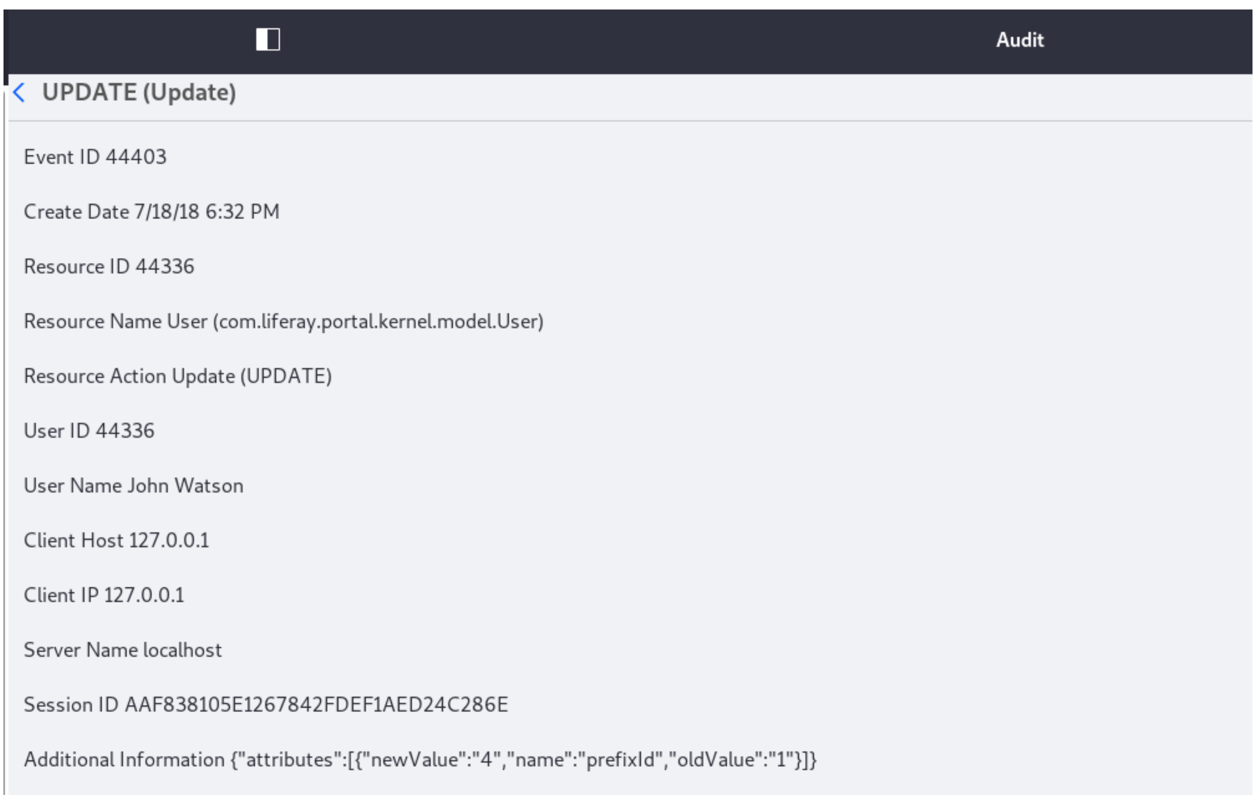

Figure 70.2: Click an event in the list to show its details. The details for this event show that John Watson updated his user account's `prefixId` from 1 to 4. The `prefixId` represents a name prefix like Dr., Mr., Mrs., or Ms.

As you can see, depending on how many users you have, this list can get populated very quickly.

That's why page view events aren't displayed by default. They'll clutter up your audit report, since they'll definitely be the most frequent event.

Note: You can add page view events to your audit report, but keep in mind that doing so adds LOTS of events. If you're a glutton for this kind of punishment, add this property to your `portal-ext.properties` file:

```
audit.message.com.liferay.portal.kernel.model.Layout.VIEW=true
```

Liferay DXP's code refers to pages as *layouts*. Setting this property to `true` therefore records audit events for page views. It's turned off by default because this is too fine-grained for most installations.

Once you've added the property, restart your server.

Finding Audit Events

Finding what you want in a big list of events is like searching for a needle in a haystack. This is why the Audit app has a robust search mechanism. By default, it looks pretty simple: there's only a single search field. Clicking the *magnifier* icon next to the search bar, however, reveals an advanced search dialog broken out by various fields you can use in your search.

Here are the available search options:

Match: Search for matches to *all* the fields you've specified or *any* single field.

User ID: The user ID to search for. This is usually the User who performed some action you'd like to audit.

User Name: The user name to search for. This is often easier than searching for a user ID, especially if you don't have access to the database containing the user ID.

Resource ID: The ID of the resource that was modified or viewed in this audit record.

Class Name: The name of the resource that was modified or viewed in this audit record. For example, you could search for user resources to see if someone modified a user's account.

Resource Action: An action performed on the resource. This could be any of these: `add`, `assign`, `delete`, `impersonate`, `login`, `login_failure`, `logout`, `unassign`, or `update`.

Session ID: The session ID to search for. You can use this to correlate a session ID from your web server logs with activity in Liferay DXP.

Client IP: The IP address of the client that performed the activity you wish to audit.

Client Host: The host name of the client that performed the activity you wish to audit.

Server Name: The name of the server in which the activity occurred. If you're using a cluster, each member of the cluster can be individually queried.

Server Port: The server port in which the activity occurred. You need this if you run a vertical cluster of multiple VMs on the same machine.

Start Date: The low end of the date range you wish to search for.

End Date: The high end of the date range you wish to search.

For example, to check if someone unassigned a User from a particular Role, you might search for a resource name of *user* and a resource action of *unassign*. Once you have the search results, you can click any of the returned records to see that record's detail page.

As you can see, the Audit app shows you what's happening as Users make changes. Use this information to troubleshoot problems, determine ownership of particular actions, or, as Harry (from the story in the introduction) is about to do, find out who made permission changes they weren't supposed to make.

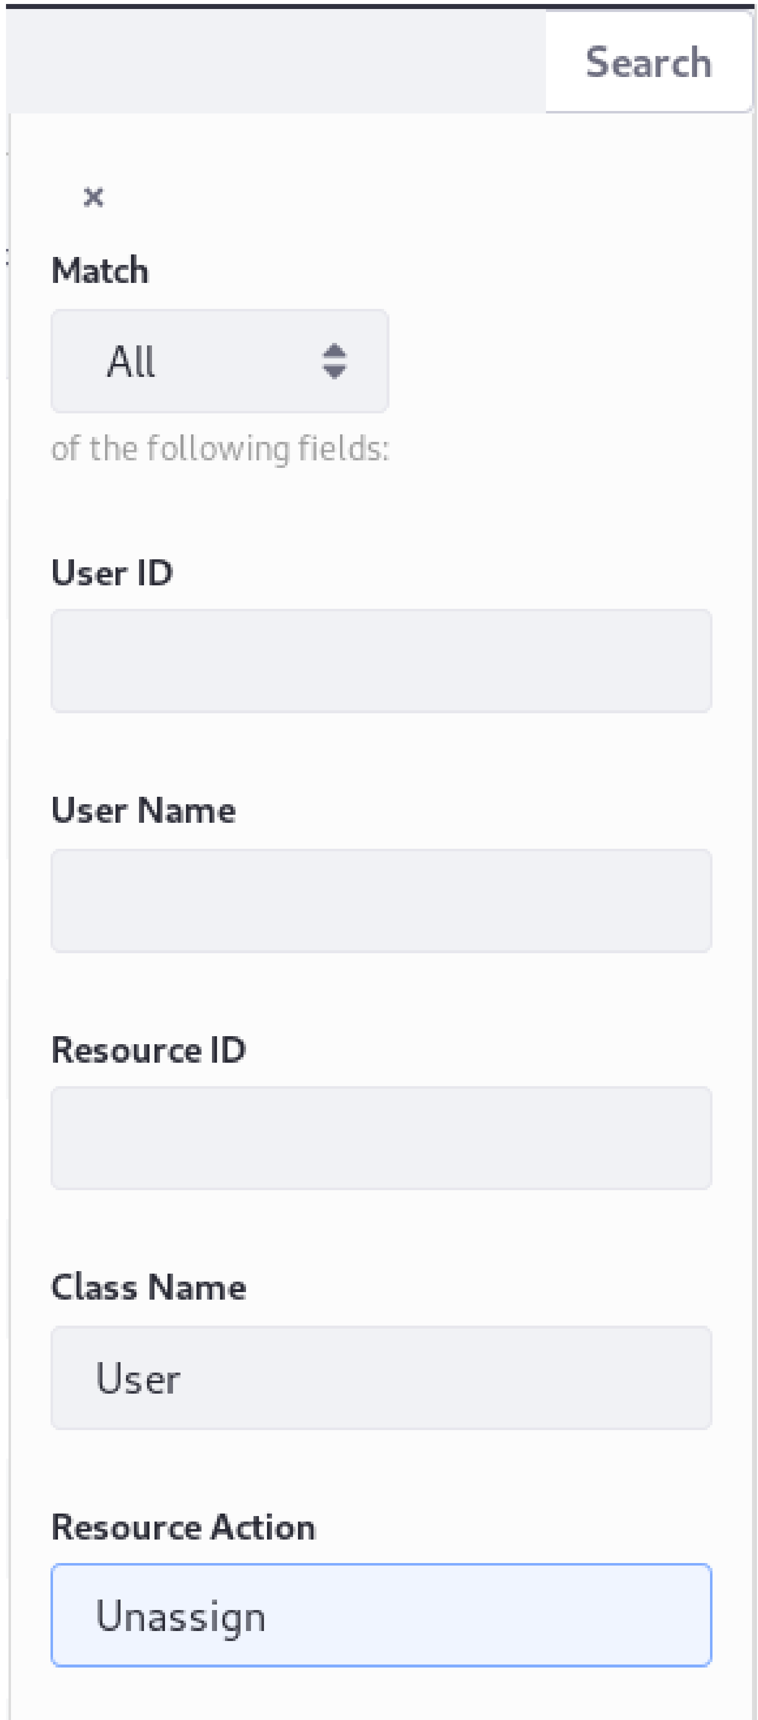

Figure 70.3: Searching for audit events is easy with the Audit app's advanced search form. You can specify various search criteria to find the types of events you want.

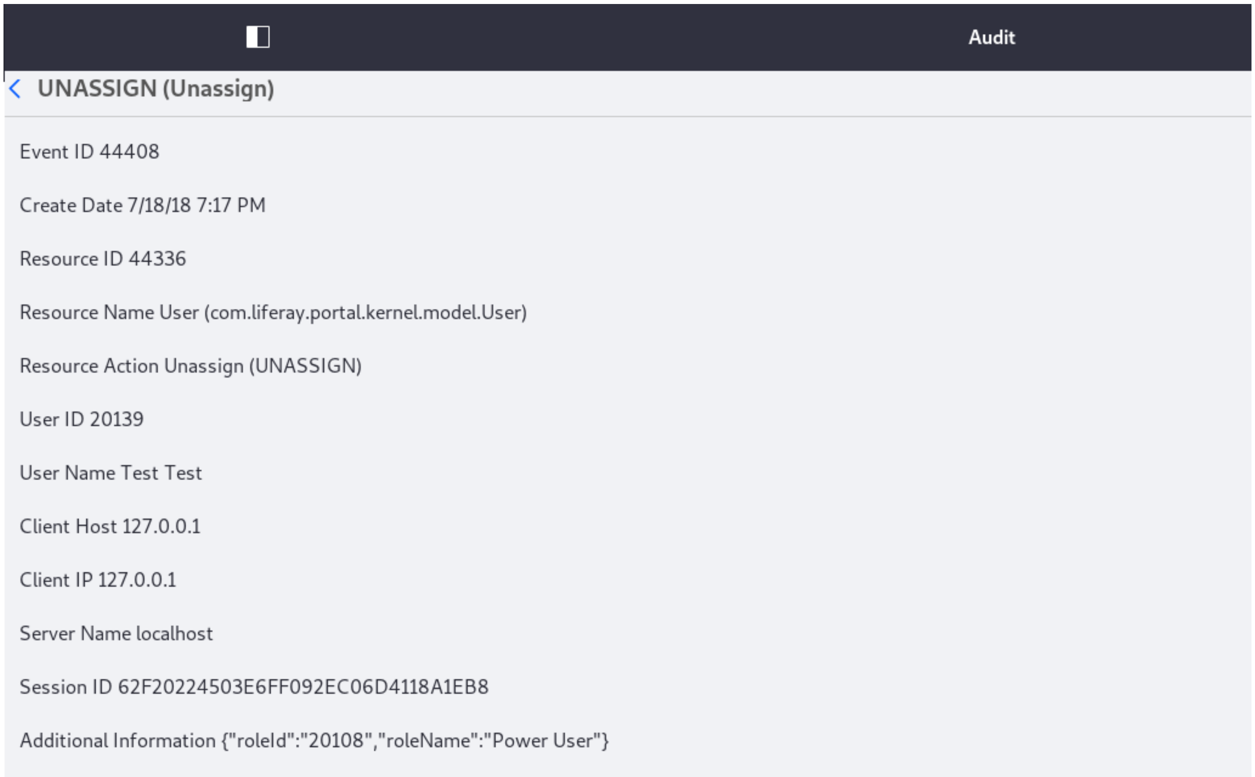

Figure 70.4: This record shows that the default administrative user removed the Power User Role from the User Test Test.

70.2 Configuring Audits

Audits are enabled by default. The Audit app reports audit events, but you can also report them in Liferay DXP's logs or console, enable them for scheduled jobs, or disable them entirely.

There are two main ways to configure Liferay DXP:

1. Edit a configuration via the Control Panel. This saves the configuration to the database.
2. Edit a configuration via an OSGi configuration file (`.config` file) in your `[LIFERAY_HOME]/osgi/configs` folder.

These methods apply to each of the audit configuration options explained below.

Reporting Audit Events in Liferay's Logs and Console

Follow these steps to use the Control Panel to configure the reporting of log events in Liferay DXP's log and console:

1. Go to *Control Panel* → *Configuration* → *System Settings*, and select *Audit* from the *Security* section.
2. In the *SYSTEM SCOPE* column on the left, select *Logging Message Audit Message Processor*.

3. Select the *Enabled* checkbox to report audit events in Liferay DXP's log.

4. Select the *Output to Console* checkbox to report audit events in the console.

5. In the *Log Message Format* selector menu, select the format for the audit events (CSV or JSON).

6. Click *Save* when you're finished.

Alternatively, you can make the same configuration via an OSGi configuration file:

1. Create a file called `com.liferay.portal.security.audit.router.configuration.LoggingAuditMessageProcessorCon`

2. Add these properties to the file:

   ```
   enabled="true"
   logMessageFormat="CSV"
   #logMessageFormat="JSON"
   outputToConsole="true"
   ```

   Note that these are the same options set in the Control Panel. Edit them as you see fit.

3. Deploy the file to the `[LIFERAY_HOME]/osgi/modules` folder. Note that the Liferay Home folder is typically the application server's parent folder.

Regardless of your configuration approach, you must also extend Liferay DXP's `log4j-ext.xml` file to configure Log4j (Liferay DXP's logging implementation) to log messages produced by the appropriate class to the appropriate file. To do so, create a `portal-log4j-ext.xml` file in `[LIFERAY_HOME]/tomcat-[version]/webapps/ROOT/WEB-INF/classes/META-INF` with this configuration:

```
<?xml version="1.0"?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">

<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">

    <!-- additional audit logging -->

    <appender name="auditFile" class="org.apache.log4j.rolling.RollingFileAppender">
        <rollingPolicy class="org.apache.log4j.rolling.TimeBasedRollingPolicy">
            <param name="FileNamePattern" value="@liferay.home@/logs/audit.%d{yyyy-MM-dd}.log" />
        </rollingPolicy>
        <layout class="org.apache.log4j.EnhancedPatternLayout">
            <param name="ConversionPattern" value="%d{ABSOLUTE} %-5p [%t][%c{1}:%L] %m%n" />
        </layout>
    </appender>

    <category name="com.liferay.portal.security.audit.router.internal.LoggingAuditMessageProcessor">
        <priority value="INFO" />
        <appender-ref ref="auditFile"/>
    </category>
</log4j:configuration>
```

This configures Log4j to record INFO level messages from the class `com.liferay.portal.security.audit.router.i` to a file called `audit.yyyy-MM-dd.log` in the `[LIFERAY_HOME]/logs` folder. Adjust the audit file properties or log level to your liking.

Configuring Audit Events for Scheduled Liferay Jobs

By default, scheduled jobs don't trigger audit events. Follow these steps to enable them via the Control Panel:

1. Go to *Control Panel* → *Configuration* → *System Settings*, and select *Infrastructure* from the *Platform* section.
2. In the *SYSTEM SCOPE* column on the left, select *Scheduler Engine Helper*.
3. Select the checkbox for *Audit scheduler job enabled* and click *Save*.

Alternatively, you can make the same configuration via an OSGi configuration file:

1. Create a file called `com.liferay.portal.scheduler.configuration.SchedulerEngineHelperConfiguration.config`.
2. Add this property to the file:

```
auditSchedulerJobEnabled=true
```

3. Deploy the file to the `[LIFERAY_HOME]/osgi/modules` folder. Note that the Liferay Home folder is typically the application server's parent folder.

Auditing scheduled jobs is a smart choice if there's a chance someone with a dubious competence level would try to schedule jobs, as you'll find out below in the conclusion of our story.

Enabling or Disabling Audit Events Entirely

Audit events are enabled by default. Follow these steps to disable them via the Control Panel:

1. Go to *Control Panel* → *Configuration* → *System Settings* and then click *Audit* in the *Security* section.
2. Uncheck the *Enabled* box. Note that when auditing is enabled, you can adjust the audit message max queue size from its default value.

Alternatively, you can enable or disable auditing via an OSGi configuration file:

1. Create a file called `com.liferay.portal.scheduler.configuration.SchedulerEngineHelperConfiguration.config`.
2. Add these properties to the file. You can adjust their values as desired:

```
enabled="true"
auditMessageMaxQueueSize="200"
```

3. Deploy the file to the `[LIFERAY_HOME]/osgi/modules` folder. Note that the Liferay Home folder is typically the application server's parent folder.

The End of the Story

"Okay," says Harry, "let's fire up the audit system and see if we can figure out what happened."

You and Dick stand behind Harry's chair and watch as he enters a query into a form on the Audit app. After clicking *search*, the screen fills up with audit events.

"Wow, that's a lot of unassign events." Harry says. "And look who the culprit is," he adds sarcastically.

"Who's Melvin Dooitrong?" Dick asks.

"That's my new intern," Harry says. "I'm gonna kill him." Harry pushes out his chair and walks down the row of cubes to the end, where a kid no more than 20 years old with disheveled hair sits, earbuds in his ears.

"Hey Melvin," Harry says as Melvin turns around to face him. "Didn't I ask you to move that set of users from Site membership to Organization membership?"

"Yeah," Melvin says, "I did that already."

"How'd you do it?"

"It was going to take a while to do it manually, so I wrote a script and executed it in the scripting host," Melvin replies, matter-of-factly.

"You did, did you? Well, guess what? Your script removed *everybody* from *all* Sites."

"*What?*"

"Yeah, and now you're going to start adding them back, one by one, manually, starting with Mr. Jones...."

Tom and Dick back away slowly from Melvin's cube as Harry and Melvin continue to have their—let's call it a discussion. One thing is clear: they're having a better day than Melvin is.

Chapter 71

Managing Apps

Apps run on the platform Liferay DXP provides. The web experience management, collaboration, and business productivity features all consist of apps. Even the Control Panel consists of configuration apps. You can also add to or change built-in functionality by installing other apps. There are several ways to manage, find, and install apps. This section covers these topics and more.

71.1 Managing and Configuring Apps

Liferay DXP is a platform for deploying apps that comprise modules and components. It has conveniences for managing apps and app management best practices for maximizing stability. Best practices in production environments involve stopping the server before applying changes, but in cases where this isn't feasible, you can "hot deploy" changes several different ways.

There are two places in the Control Panel where you can manage and configure apps: the App Manager and the Components listing. The App Manager manages apps in the OSGi framework. You can use the App Manager to install, activate, deactivate, and delete apps. You can manage apps at the app and module levels.

The Components listing views and manages apps at the OSGi component level. It differs from the App Manager by showing apps by type (portlet, theme, and layout template), and setting app permissions. You can use the Components listing to activate and deactivate apps, but it can't install or delete apps.

Start with learning app management best practices in production, or wherever you want to maximize stability.

Managing Apps in Production

Not all apps are designed to be "hot deployed"—deployed while the server is running. Deploying that way can cause instabilities, such as class loading leaks and memory leaks. On production systems, avoid "hot deploying" apps and configurations whenever possible.

If you're installing an app or a component configuration on a production system and stopping the server is feasible, follow these steps:

1. Stop your server.

2. Copy your app (`.lpkg`, module `.jar`, or plugin `.war`) to your `[Liferay Home]/deploy` folder, or copy your component configuration (`.config` file) to the `[Liferay Home]/osgi/configs` folder. The Liferay Home folder is typically the app server's parent folder.

3. Start your server.

If you're running in cluster, follow the instructions for updating a cluster.

Warning: Avoid repeatedly "hot deploying" new versions of apps that aren't designed for "hot deployment".

If it's not feasible to stop your server or you're app *is* designed for "hot deployment", Liferay DXP provides several "hot deployment" conveniences. Except where stopping/starting the server is explicitly mentioned, the practices described in the rest of this article and in the following articles involve "hot deployment".

Using the App Manager

Access the App Manager by selecting *Control Panel* → *Apps* → *App Manager*. The App Manager lists your apps. The *Filter and Order* menu lets you filter and order by category, status, or title. Click the up or down arrows to perform an ascending or descending sort, respectively. To search for an app or module, use the search bar. This is often the quickest way to find something.

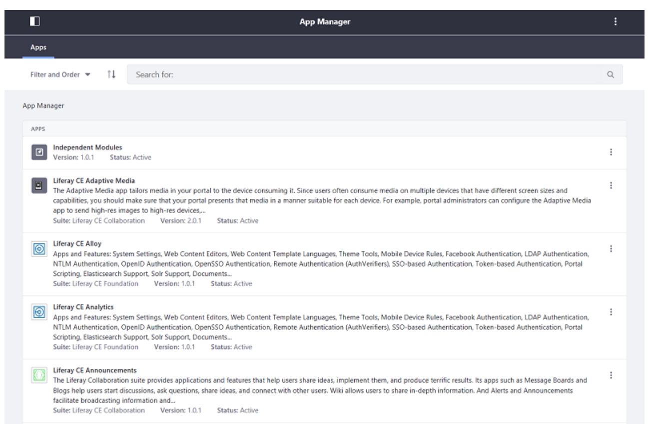

Figure 71.1: The App Manager lets you manage the apps installed in your Liferay DXP instance.

Each item listed in the table contains a description (if available), version, and status. Here are the statuses:

- **Installed:** The item is installed to Liferay DXP.
- **Resolved:** The item's dependencies are active. Resolved items can typically be activated. Some items, however, can't be activated and are intended to remain in the Resolved state (e.g., WSDD modules containing SOAP web services).
- **Active:** The item is running in Liferay DXP.

Clicking each item's Actions button ( ) brings up a menu that lets you activate, deactivate, or uninstall that item.

To view an item's contents, click its name in the table. If you click an app, the app's modules are listed. If you click a module, the module's components and portlets appear. The component level is as far down as you can go without getting into the source code. At any level in the App Manager, a link trail appears that lets you navigate back in the hierarchy.

For information on using the App Manager to install an app, see the article Installing Apps Manually.

Next, you'll learn how to use the Components listing.

Using the Components Listing

Access the Components listing by selecting *Control Panel* → *Configuration* → *Components*. The Components listing first shows a table containing a list of installed portlets. Select the type of component to view—portlets, themes, or layout templates—by clicking the matching tab on top of the table. To configure a component, select its name in the table or select *Edit* from its Actions button ( ). Doing either opens the same configuration screen.

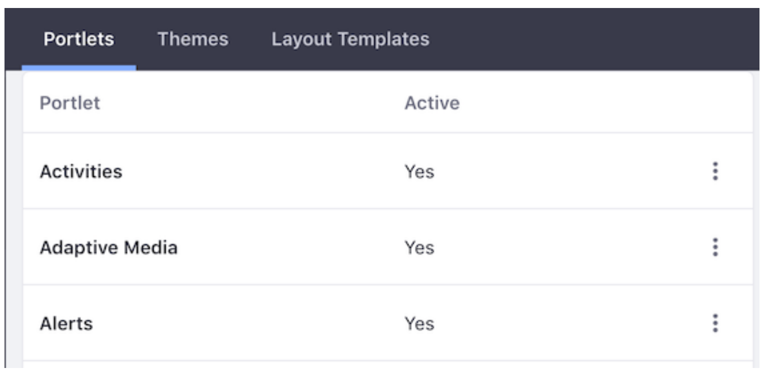

Figure 71.2: The Components listing lets you manage the portlets, themes, and layout templates installed in your Liferay DXP instance.

The configuration screen lets you view a component's module ID and plugin ID, activate or deactivate the component, and change the component's Add to Page permission. The component's module ID and plugin ID appear at the top of the screen. You can activate or deactivate a component by checking or unchecking the *Active* checkbox, respectively. To change a component's Add to Page permission for a role, select the role's *Change* button in the permissions table. This takes you to *Control Panel* → *Users* → *Roles*, where you can change the component's permissions for the selected role.

71.2 Using the Liferay Marketplace

Liferay Marketplace is a hub for sharing, browsing, and downloading apps. Marketplace leverages the entire Liferay ecosystem to release and share apps in a user-friendly, one-stop shop.

There are two ways to access the Marketplace.

1. **Via the website:** Using your favorite browser, navigate to the Marketplace at web.liferay.com/marketplace. If you're new to Marketplace, this is the easiest way to access it. You can browse Marketplace without signing in with your liferay.com account.

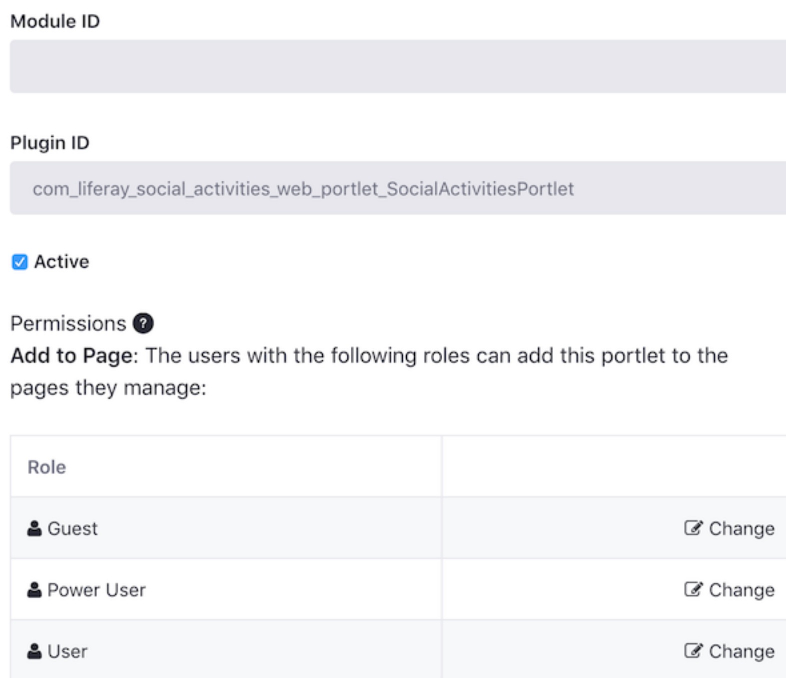

Figure 71.3: You can activate or deactivate a component, and change its permissions.

2. **Via the Control Panel:** In the Control Panel, navigate to *Apps* → *Store*. To view Marketplace, you must sign in with your liferay.com account.

No matter how you access Marketplace, you'll see the same content. Note that to download apps, you must have a liferay.com account and agree to the Marketplace Terms of Use.

Here you'll learn how to,

- Find and purchase apps
- Manage purchased apps
- Renew purchased apps

Start with finding and purchasing the apps you want.

Finding and Purchasing Apps

If you've used an app store before, Marketplace should be familiar. You'll see apps in the center of the page, in the following sections:

- Featured Apps: Liferay features a different set of apps each month.
- New and Interesting: The latest apps added to Marketplace.
- Most Viewed in the Past Month: The top 5 most viewed apps in the last month.
- Themes and Site Templates: Apps that change your Liferay instance's look and feel.

- App categories: Communication, productivity, security, etc.

The *See All* button shows more info in all these sections. At the top of the page, you can search Marketplace by category, Liferay version, and price. To browse by category, click the *Categories* menu at the top-right of the page.

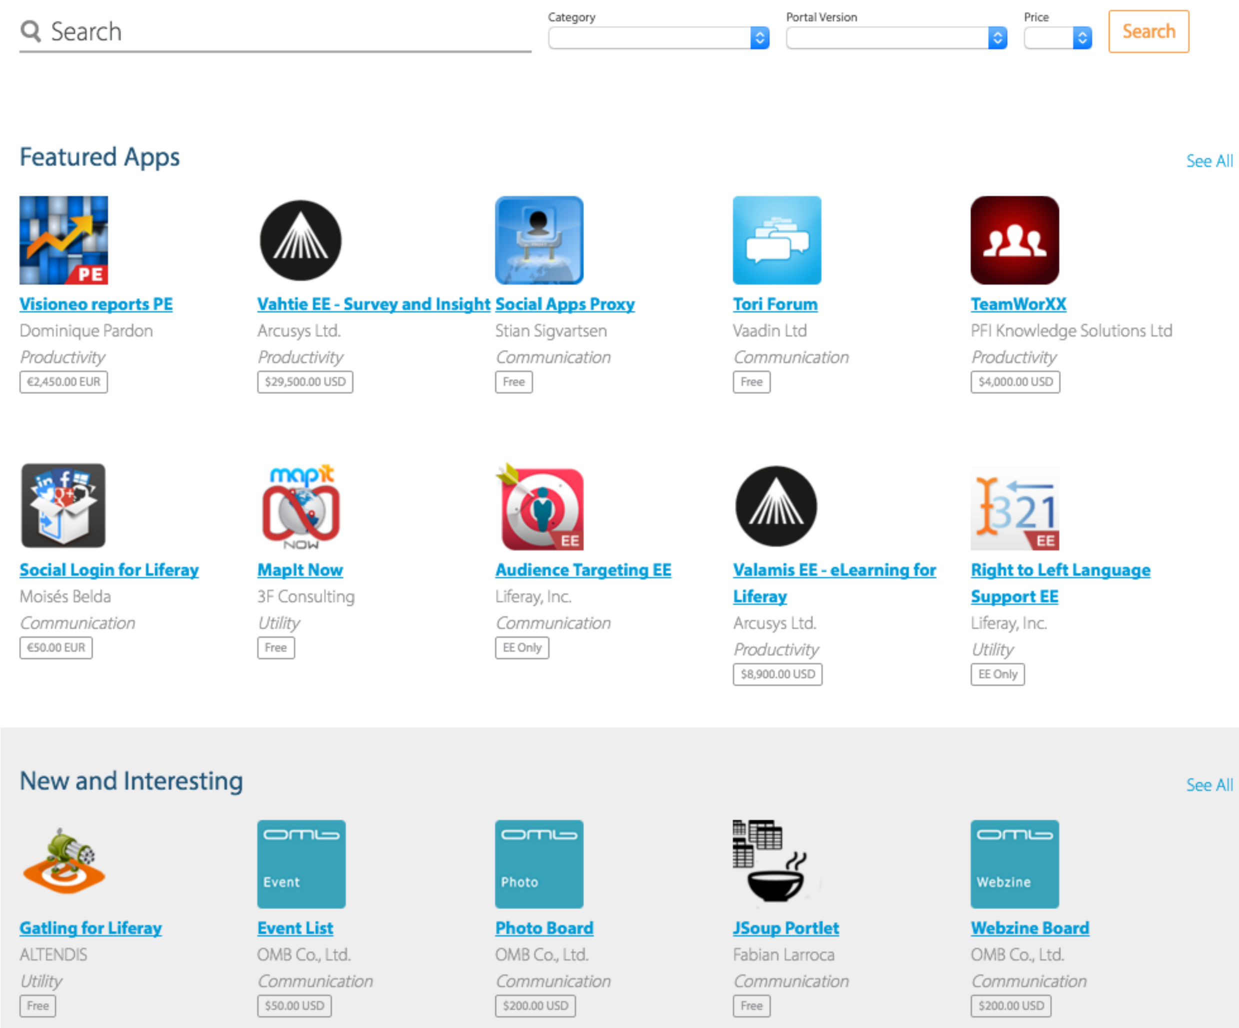

Figure 71.4: The Liferay Marketplace home page lets you browse and search for apps.

Click an app to view its details. This includes its description, screenshots, price, latest version, number of downloads, a link to the developer's website, a link to the app's license agreement, and a purchase button (labeled Free or Buy, depending on the price). You can also view the app's version history, read reviews left by other users, or write your own review.

The purchase button prompts you to choose a purchase type. You can purchase an app for your personal account, or for a Liferay project associated with your company. If you have the necessary permissions, you can also create a new project for your company. Once you select a purchase type, accept the EULA and Terms of Service, and click *Purchase*.

Once you purchase an app, you can download and install it.

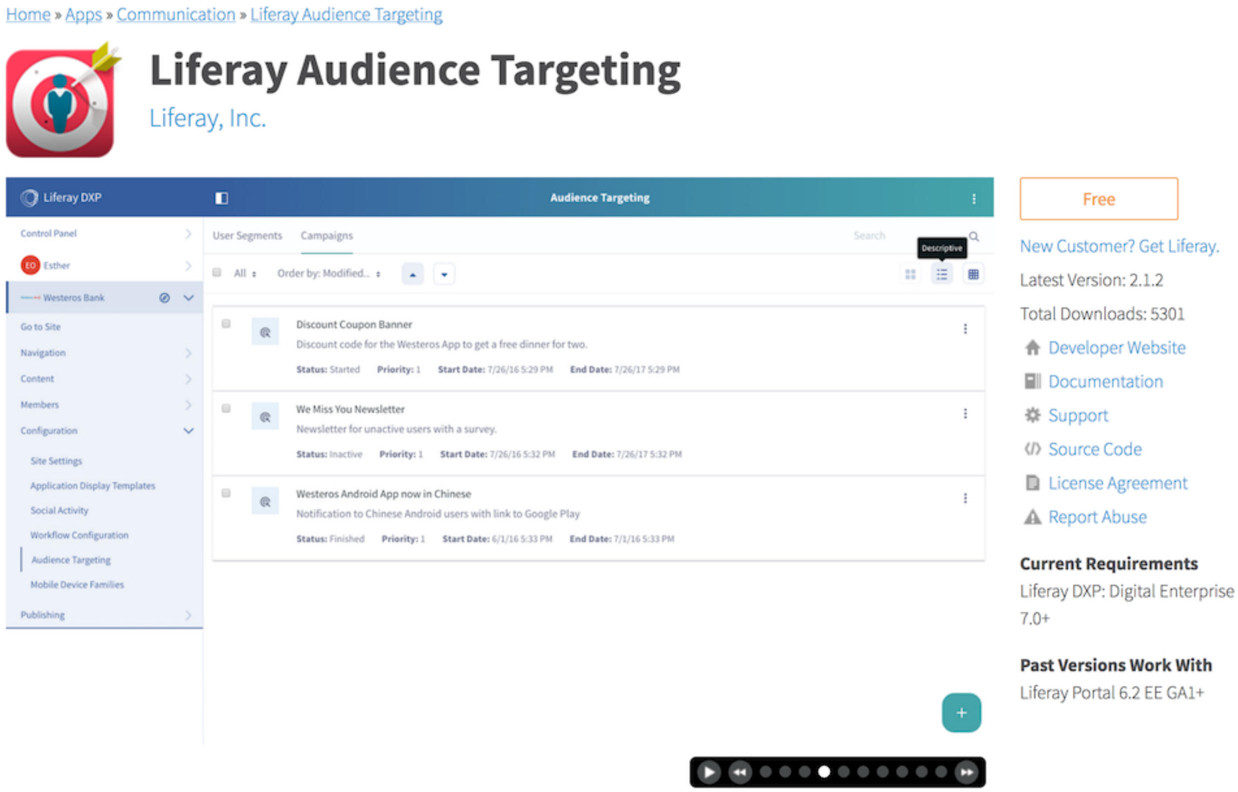

Figure 71.5: Click an app to view its details.

Warning: Not all apps are designed to be "hot deployed"—deployed while the server is running. Deploying that way can cause instabilities, such as class loading leaks and memory leaks. On production systems, avoid "hot deploying" apps whenever possible. See the best practices for managing apps in production.

An app downloads and installs immediately if you purchase it from the Control Panel. If you purchase the app on the Marketplace website, however, your receipt is displayed immediately after purchase. To download the app, click the *See Purchased* button on the bottom of the receipt, and then click the *App* button to start the download. You must then install the app manually. Alternatively, you can use Marketplace from the Control Panel to download and install the app after purchase on the Marketplace website. The next section shows you how to do this.

Note that sometimes administrators disable automatic app installations so they can manage installations manually. In this case, Marketplace apps downloaded from the Control Panel are placed in the `deploy` folder in Liferay Home. Administrators must then manually install the app from this folder. Manual install is also required if the server is behind a corporate firewall or otherwise lacks direct Marketplace access. Regardless of how the app is downloaded, the manual install process is the same. For details, see the article Installing Apps Manually.

Managing Purchased Apps

Important: When uninstalling an app or module, make sure to use the same agent you used to install the app. For example, if you installed it with Marketplace, uninstall it with Marketplace. If you installed it with the file system, use the file system to uninstall it. If you installed it with the App Manager, however, use Blacklisting to uninstall it.

There are two places to manage your purchased apps:

1. Your liferay.com account's home page. After signing in, click the user menu at the top-right and select *Account Home*. Note that your home page is distinct from your profile page. Your home page is private, while your profile page is public. On your home page, select *Apps* from the menu on the left to view your projects. Select a project to view its registered apps. Clicking an app lets you view its versions. You can download the version of the app that you need. This is especially useful if you need a previous version of the app, or can't download the app from the Control Panel.

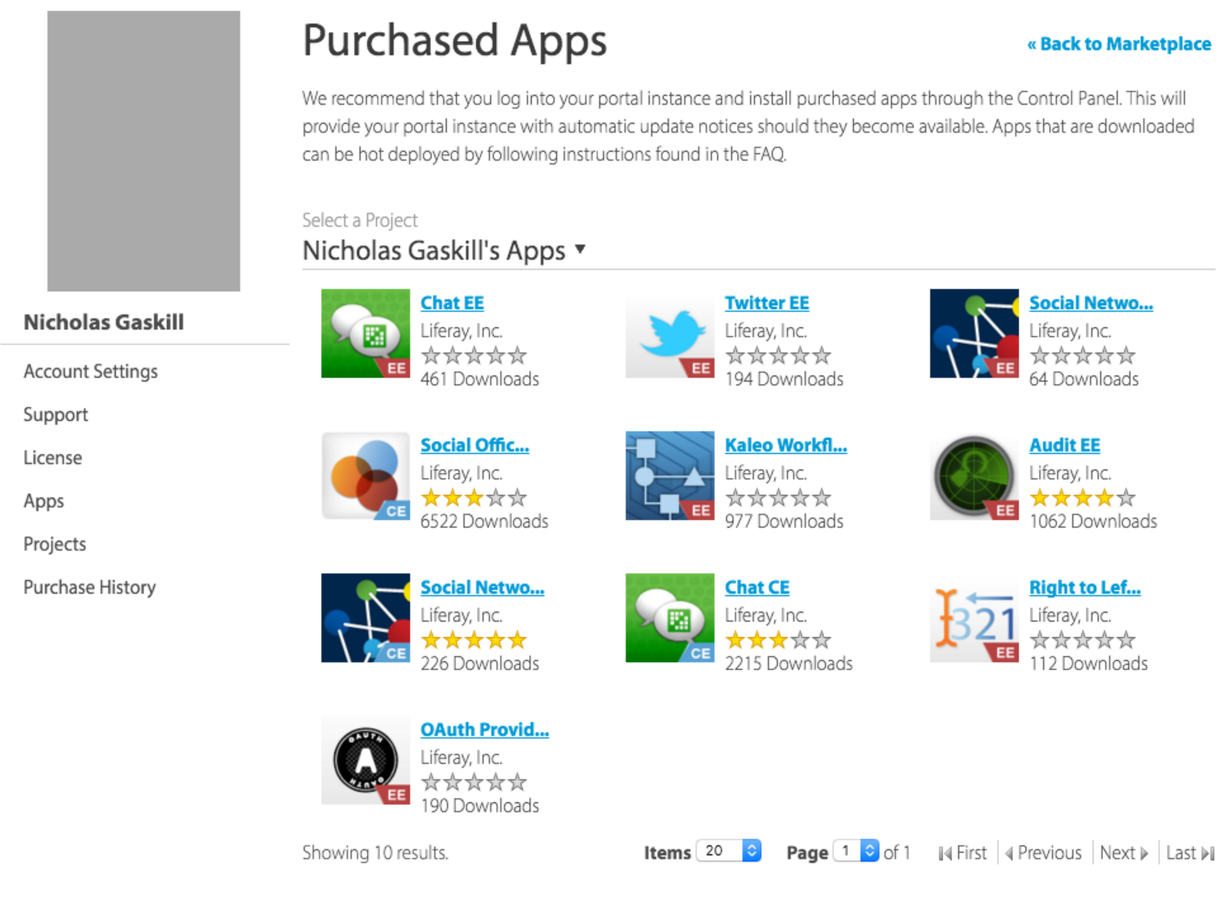

Figure 71.6: You can manage your purchased apps from your liferay.com account's home page.

2. From the Control Panel. Navigate to *Apps* → *Purchased* to see your purchased apps. A button next to each app lets you install or uninstall the app. If the app isn't compatible with your

Liferay DXP version, *Not Compatible* is displayed in place of the button. Additional compatibility notes are also shown, such as whether a newer version of the app is available. You can also search for an app here by project, category, and title. Clicking the app takes you to its Marketplace entry.

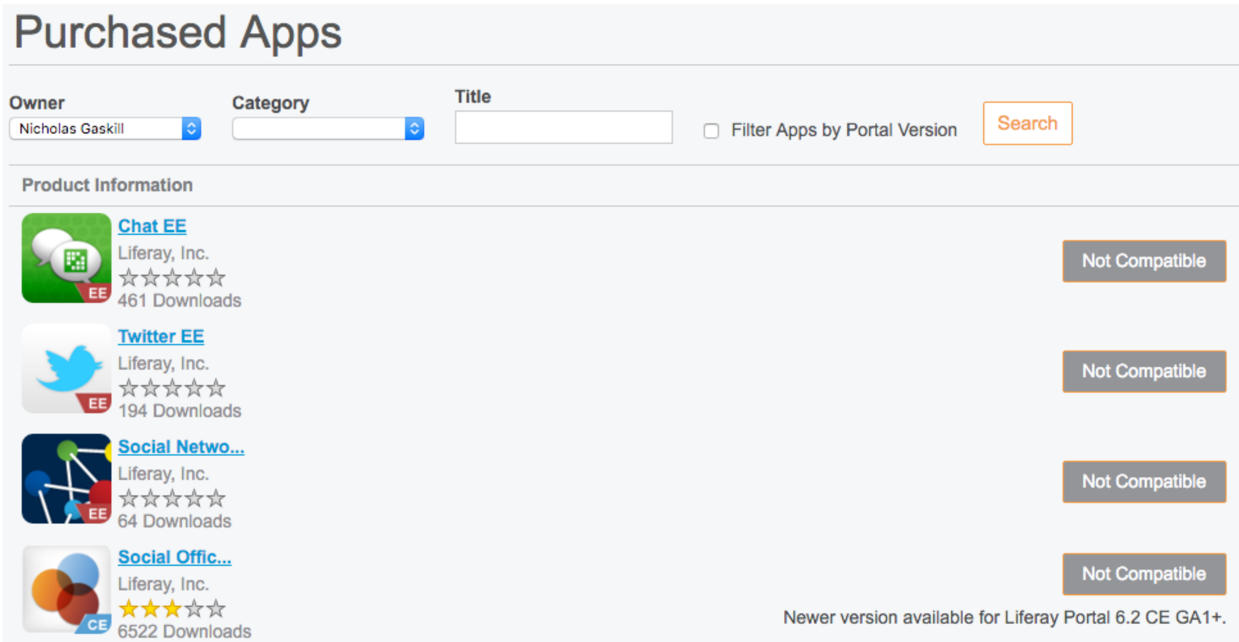

Figure 71.7: You can also manage your purchased apps from within a running Liferay instance.

Renewing a Purchased App

To continue using a purchased app whose license terms are non-perpetual, you must renew your app subscription, register your server to use the app, and generate a new activation key to use on your server. Here are the steps:

1. Go to https://web.liferay.com/marketplace.
2. Click your profile picture in the upper right corner and select *Purchased Apps*. The Purchased Apps page appears and shows your app icons organized by project.
3. Click your app's icon. Your app's details page appears.
4. Click *Manage Licenses*.
5. Select *Register New Server*.
6. Select the most recent *Order ID* (typically the order that has no registered servers).
7. Fill in your server's details.
8. Click *Register*.
9. Click *Download*. The new app activation key to use on your server downloads.

10. Copy the activation key file to your `[Liferay Home]/deploy` folder.

    You can continue using the application on your server.

71.3 Installing Apps Manually

By default, apps you download from the Control Panel via Liferay Marketplace install automatically. But what if the app you want to install isn't on Marketplace? What if all you have is the app's file? In this case, you must install the app manually. This article shows you how to do this. You can use the procedures here to install any app.

Warning: Not all apps are designed to be "hot deployed"—deployed while the server is running. Deploying that way can cause instabilities, such as class loading leaks and memory leaks. On production systems, avoid "hot deploying" apps whenever possible. See the best practices for managing apps in production.

Important: When uninstalling an app or module, make sure to use the same agent you used to install the app. For example, if you installed it with Marketplace, uninstall it with Marketplace. If you installed it with the file system, use the file system to uninstall it. If you installed it with the App Manager, however, use Blacklisting to uninstall it.

Using the Control Panel to Install Apps

To install an app manually from the Control Panel, navigate to *Control Panel* → *Apps* → *App Manager*, and select *Upload* from the options button ( ). In the Upload dialog, choose the app on your machine and then click *Install*. When the install completes, close the dialog and you're ready to roll!

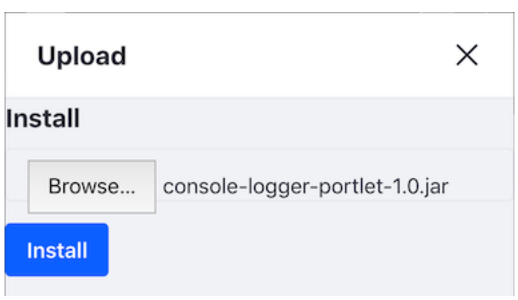

Figure 71.8: Installing an app via the Upload dialog in the App Manager is very simple.

Using Your File System to Install Apps

To install an app manually on the Liferay DXP server, put the app in the `[Liferay Home]/deploy` folder (the Liferay Home folder is typically the app server's parent folder). That's it. The auto deploy mechanism takes care of the rest.

You might now be thinking, "Whoa there! What do you mean by 'the rest?' What exactly happens here? And what if my app server doesn't support auto deploy?" These are fantastic questions! When you put an app in the `[Liferay Home]/deploy` folder, the OSGi container deploys the app to the appropriate subfolder in `[Liferay Home]/osgi`. By default, the following subfolders are used for apps matching the indicated file type:

- `marketplace`: Marketplace LPKG packages
- `modules`: OSGi modules
- `war`: WAR files

You can, however, change these subfolders by setting the properties `module.framework.base.dir` and `module.framework.auto.deploy.dirs` in a `portal-ext.properties` file. These properties define the `[Liferay Home]/osgi` folder and its hot deploy subfolders, respectively. The default settings for these properties in the `portal.properties` file are as follows:

```
module.framework.base.dir=${liferay.home}/osgi

module.framework.auto.deploy.dirs=\
    ${module.framework.base.dir}/configs,\
    ${module.framework.base.dir}/marketplace,\
    ${module.framework.base.dir}/modules,\
    ${module.framework.base.dir}/war
```

Note that the `configs` subfolder isn't for apps: it's for configuration files imported from other Liferay DXP instances.

But what happens if your app server doesn't support auto deploy? No problem! Liferay DXP's module framework (OSGi) enables auto deploy. Any app server running Liferay DXP therefore also supports this auto deploy mechanism.

Manually Deploying an LPKG App

When manually installing an LPKG app, the installation may hang with a server log message like this:

```
14:00:15,789 INFO  [com.liferay.portal.kernel.deploy.auto.AutoDeployScanner][AutoDeployDir:252] Processing Liferay Push 2.1.0.lpkg
```

This happens when LPKG apps have the `restart-required=true` property in their `liferay-marketplace.properties` file (inside the LPKG file). This property setting specifies that a server restart is required to complete the installation.

71.4 App Types

There are several different kinds of apps. Some apps can even contain other apps. The types of apps you can install include:

- OSGi Modules
- Portlets
- Web Plugins
- Templates
- Themes

Read on to learn about these app types.

OSGi Modules

Since Liferay DXP runs on OSGi, apps can be implemented as OSGi modules. An OSGi module is a JAR file adapted to run on OSGi. Although it's possible for a single module to implement a single app, an app typically consists of multiple modules that are packaged together. Also note that apps in OSGi modules aren't required to have a UI. For example, Liferay DXP can run OSGi modules that expand built-in APIs without requiring any user interaction. This is crucial for developers that must leverage custom APIs. By providing such an API via one or more OSGi modules, you can let developers leverage your API.

OSGi modules can also contain apps that have a UI: portlets. The next section discusses these.

Portlets

Portlets are small web applications that run in a portion of a web page. For example, the built-in Blogs app is a portlet. Portlet applications, like servlet applications, are a Java standard implemented by various portal server vendors. The JSR-168 standard defines the portlet 1.0 specification, the JSR-286 standard defines the portlet 2.0 specification, and the JSR-362 standard defines the portlet 3.0 specification. A Java standard portlet should be deployable on any portlet container that supports the standard. Portlets are placed on the page in a certain order by the end user and are served up dynamically by the portal server. This means certain things that apply to servlet-based projects, such as control over URLs or access to the `HttpServletRequest` object, don't apply in portlet projects because the portal server generates these objects dynamically.

Portlets can be composed of OSGi modules (recommended), or contained in WAR files. For information on developing OSGi modules for Liferay, including portlets, see Liferay's developer tutorials.

Web Plugins

Web plugins are regular Java EE web modules designed to work with Liferay DXP. You can integrate with various Enterprise Service Bus (ESB) implementations, as well as Single Sign-On implementations, workflow engines, and so on. These are implemented as web modules used by Liferay DXP portlets to provide functionality.

Templates and Themes

Templates and themes are plugins that change Liferay DXP's appearance. Templates (layout templates) control how you can arrange portlets on a page. They make up a page's body (the large area into which you can drag and drop portlets). There are several built-in layout templates. If you have a complex page layout (especially for your home page), you may wish to create a custom layout template of your own.

Themes can completely transform Liferay DXP's look and feel. Most organizations have their own look and feel standards that apply to all of their web sites and applications. By using a theme plugin, an organization can apply these standards on Liferay DXP. There are many available theme plugins on Liferay's web site and more are being added every day. This makes it easy for theme developers, as they can customize existing themes instead of writing a new one from scratch.

For information on developing themes and templates, see this section of tutorials.

Liferay Marketplace App Packages

Regardless of app type, each Liferay Marketplace app is distributed in an LPKG package. The LPKG package contains Marketplace metadata and the files the app needs to run. Note that it's possible for an LPKG package to contain multiple apps. For example, a single LPKG package can contain several portlets. This is common in cases where an app requires a Control Panel portlet for administrators, and another portlet for end users.

71.5 Blacklisting OSGi Modules and Components

Blacklists are used for good and evil. An evil blacklist penalizes unfairly; a good blacklist protects. Liferay DXP's OSGi module and component blacklists are files that prevent particular modules from installing and particular components from enabling. This saves you the trouble of uninstalling and disabling them individually with the Application Manager, Components list, or Gogo shell.

Blacklisting Modules

Liferay DXP removes any installed modules on the blacklist. Blacklisted modules therefore can't be installed. The log reports each module uninstallation.

Follow these steps to blacklist modules:

1. In the Control Panel, navigate to *Configuration* → *System Settings* → *Module Container*.

2. In the Bundle Blacklist screen, add the bundle symbolic names (see the table below) for the OSGi module JARs, LPKGs, or WARs to uninstall. Click the *Save* button when you're finished.

Figure 71.9: This blacklist uninstalls the `com.liferay.docs.greeting.api` OSGi module, Liferay Marketplace LPKG, and `classic-theme` WAR.

3. To export the blacklist, click its Actions button (⋮) and then click *Export*. The blacklist configuration file then downloads (`com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.con`
   Here are contents from an example file:

```
blacklistBundleSymbolicNames=["com.liferay.docs.greeting.api","Liferay\ Marketplace","classic-theme"]
```

4. Add the bundle symbolic names of any modules not already listed that you want to prevent from installing.

   Important: Configuration values can't contain extra spaces. Extra spaces can short-circuit lists or invalidate the configuration entry.

5. To deploy the configuration file, copy it into the folder `[Liferay_Home]/osgi/configs`. The Liferay Home folder is typically the app server's parent folder.

Note: Blacklisting an LPKG uninstalls all of its internal modules.

Blacklist Bundle Symbolic Names

Type	Bundle Symbolic Name
Module JAR	`Bundle-SymbolicName` in `bnd.bnd` or `MANIFEST.MF` file
LPKG	LPKG file name without the `.lpkg` extension
WAR	Servlet context name in `liferay-plugin-package.properties` file or the WAR file name (minus `.war`), if there is no servlet context name property

Reinstalling Blacklisted Modules

To reinstall and permit installation of blacklisted modules, follow these steps:

1. Open the configuration file `com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.con`
2. Remove the symbolic names of the modules, LPKGs, or WARs from the `blacklistBundleSymbolicNames` list and save the file.

To reinstall *all* the blacklisted modules execute one of these options:

- Remove the configuration file.
- Uninstall the module `com.liferay.portal.bundle.blacklist` using the Application Manager or Felix Gogo Shell.

Note: To temporarily reinstall a module that's been blacklisted, you can remove its symbolic name from the Bundle Blacklist module in System Settings and click the *Update* button. If you want the module to install on subsequent server startup, make sure to remove the module's symbolic name from any existing module blacklist configuration file in the `[Liferay_Home]/osgi/configs` folder.

The log reports each module installation.

Blacklisting Components

Follow these steps to blacklist components:

1. In the Control Panel, navigate to *Configuration* → *System Settings* → *Module Container*.

2. In the Component Blacklist Configuration screen, add the names of components to disable and click the *Save* button.

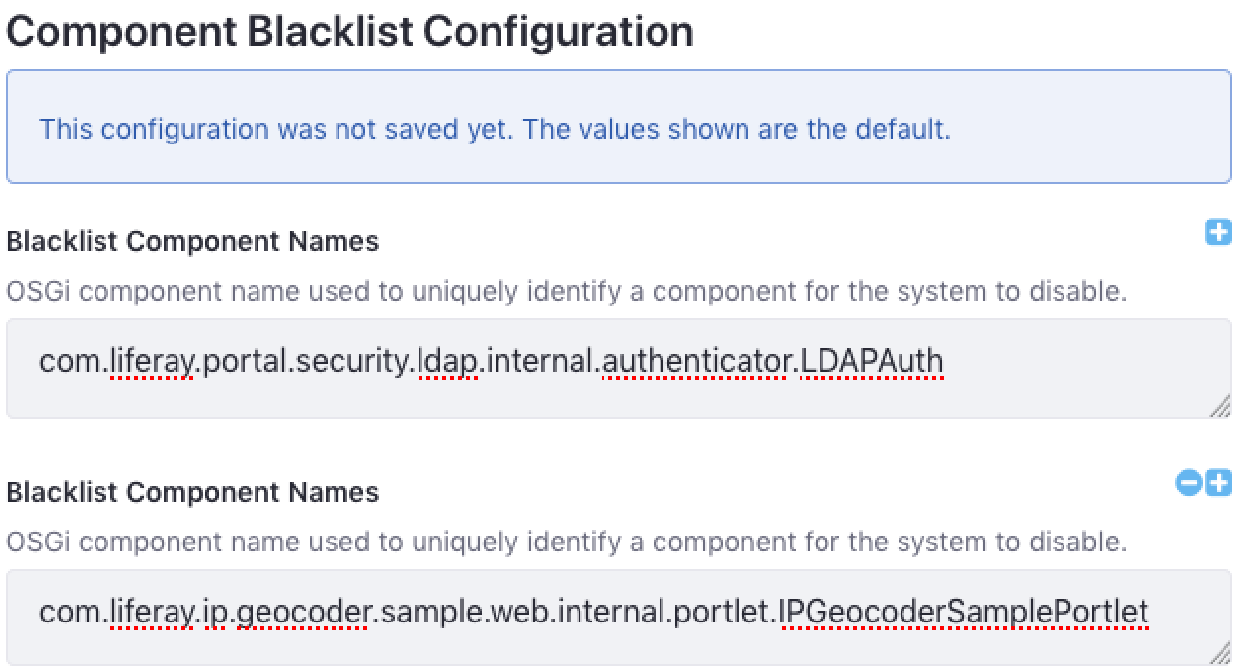

Figure 71.10: This blacklist disables the components `com.liferay.portal.security.ldap.internal.authenticator.LDAPAuth` and `com.liferay.ip.geocoder.sample.web.internal.portlet.IPGeocoderSamplePortlet`.

3. To export the blacklist, click on the Component Blacklist Configuration module's Actions button (⋮) and then click *Export*. The blacklist configuration file then downloads (`com.liferay.portal.component.blacklist.internal.ComponentBlacklistConfiguration.config`). Here are contents from an example file:

```
blacklistComponentNames=["com.liferay.portal.security.ldap.internal.authenticator.LDAPAuth","com.liferay.ip.geocoder.sample.web.internal.port
```

4. Add the names of any components not already listed (e.g., components of modules not yet installed) that you want to prevent from activating.

   Important: Configuration values can't contain extra spaces. Extra spaces can short-circuit lists or invalidate the configuration entry.

5. To deploy the configuration file, copy it into the folder `[Liferay_Home]/osgi/configs`. The Liferay Home folder is typically the app server's parent folder.

Re-enabling Blacklisted Components

To re-enable and permit enabling of blacklisted components, follow these steps:

1. Open the configuration file `com.liferay.portal.component.blacklist.internal.ComponentBlacklistConfigurati`
2. Remove the names of the components from the `blacklistComponentNames` list and save the file.

To enable *all* the blacklisted components, remove the configuration file.

Note: To temporarily reactivate a blacklisted component, remove its name from the Component Blacklist Configuration module in System Settings and click *Update*. If you want the component to activate on subsequent server startup, make sure to remove the component's name from any existing component blacklist configuration file in the `[Liferay_Home]/osgi/configs` folder.

Chapter 72

Setting Up

If Liferay DXP is anything, it's configurable. As the core is shrinking due to its increased modularity, it's important that all the applications in Liferay are also configurable.

Breaking it down, three types of applications must be configurable:

1. The platform itself
2. Liferay's out-of-the-box applications
3. Custom applications

To this end, Liferay's engineers have made the platform and its applications configurable, and created a mechanism for developers to make their applications configurable.

72.1 Where Configuration Happens

Liferay's configuration takes place in the following places:

User Interface: configuration through Liferay's UI is stored in a database. The values set in the UI always override configurations set in properties files.

Properties files: properties files that set default behavior may be included in the platform or the modules. Keep in mind that these settings can always be overridden by a system administrator in the UI. To find what properties are configurable this way, visit docs.liferay.com.

The UI location where these configuration options appear depends on the scope you want to affect with the settings you choose.

72.2 Configuration Scope

Depending on the configuration scope of a setting you change, you'll impact the platform and its applications with more or less granularity. At one end of the spectrum, you can affect the scope of the whole system. Configurations made at the system scope affect all virtual instances, Sites, and widgets. At the opposite end of the spectrum, configurations made at the widget level provide configuration settings only for that instance of the widget.

Take Language settings, for example. You can set the default language used by the virtual instance. You can also set the default language of a Site. Some applications even let you set the default language used by a specific piece of content.

Here's an overview of the available configuration scopes:

System: Configuring Liferay and its applications through System Settings provides system scoped configurations and sets default values for all virtual instances, sites, or widgets.

Virtual Instance: Configuring in Instance Settings provides settings that act on the specific virtual instance for which they are made, including Sites and widgets in the virtual instance.

Site: Configurations made at the Site scope, where you select the Site to configure in the Site selector, provide settings that take place only in that Site. Alternate configurations can be made in different Sites.

Widget Scope: Configuring a specific widget only provides a configuration for that particular widget.

Scopes in Liferay are hierarchical so that one scope can set the default values for the underlying sub-scopes. For example, making a system-wide configuration sets the default values for all virtual instances, sites or widgets of the system. If a different configuration is set at a level with more granularity (for example, the widget scope), it takes precedence over the settings made at less granular scopes (for example, the virtual instance scope).

This section contains articles on configuring Liferay at the System and Instance scopes:

System wide configuration:

- System Settings is the primary location for system configuration.
- Server Administration contains some lower-level server configuration options, such as logging.

Setting up a virtual instance:

- Virtual Instances is where virtual instances are added and edited.
- Instance Settings is the primary location for a virtual instance's configuration.
- Custom Fields is where additional database fields are added to existing virtual instance entities.

All of these are accessed through the Control Panel. Start by learning to configure modules system-wide in the System Settings Control Panel app.

Chapter 73

System Wide Settings

It can be hard to keep track of all the configuration interfaces. The Control Panel's Configuration section houses a lot of the higher level (for example, system and instance scoped) configuration options. This section considers the configuration options dealing with the *System* scope. Configuration at the system level affects all the *Virtual Instances* of Liferay in the system. If you're not sure what a Virtual Instance is, read the section on Setting Up a Liferay Instance.

There are several Control Panel sections with system scoped configurations:

- System Settings is the primary location for system configuration.
- Server Administration contains some lower-level server configuration options, such as logging.

Get started by learning about System Settings.

73.1 System Settings

Liferay DXP is modular, meaning it's composed of many applications divided into even smaller "chunks" of functionality. The applications, and sometimes even code chunks, are configurable at several scopes, as discussed in the introductory article for this section.

In System Settings, administrators make system scoped configuration changes and set system-wide default configurations for other scopes. System Settings is located in Control Panel → Configuration → System Settings.

Editing System Configurations

System Settings is organized into sections (for example, Content) and categories (for example, Workflow) based on the functionality being configured. There's also a Search bar to make finding configuration entries easier. Search for the name of a specific configuration entry, or even a specific field within an entry.

Changing a configuration isn't difficult:

1. Find the configuration you need, either by searching or browsing the sections and categories.

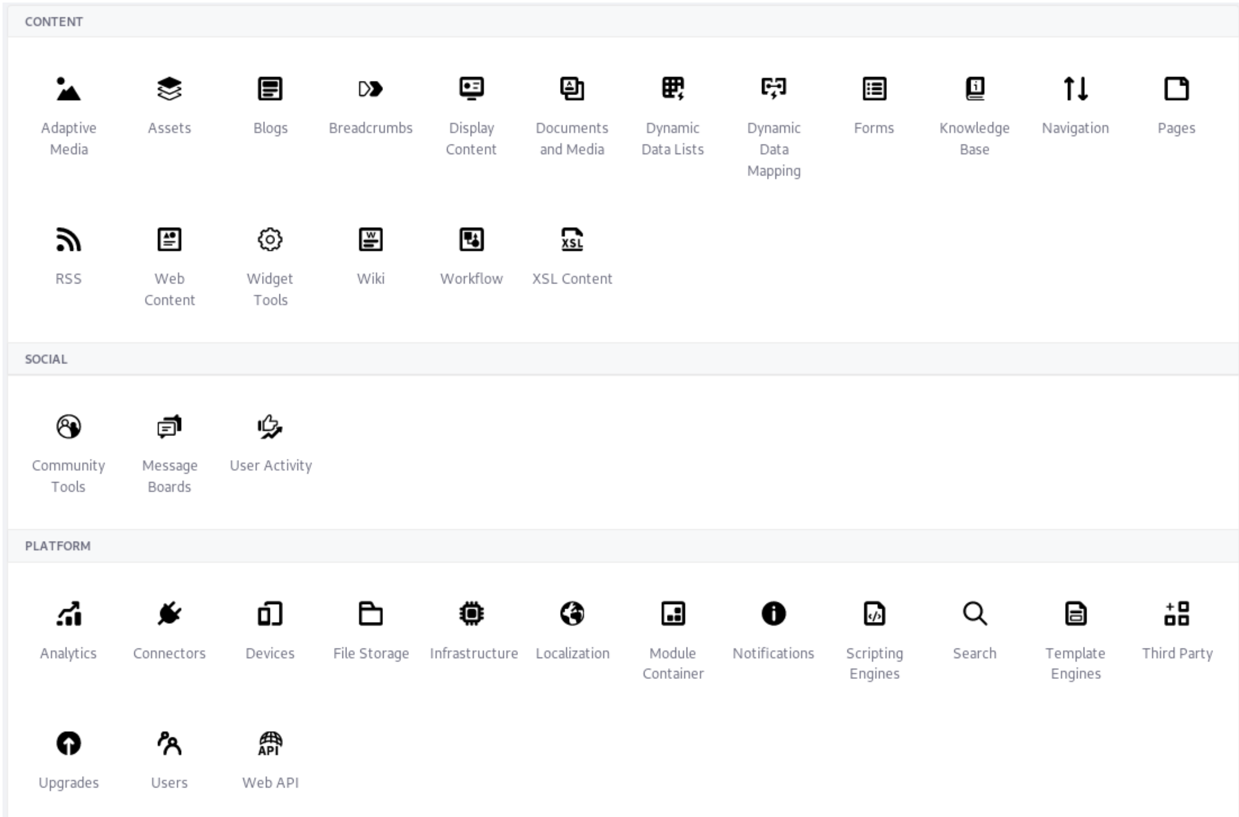

Figure 73.1: System Settings are accessed through the Control Panel.

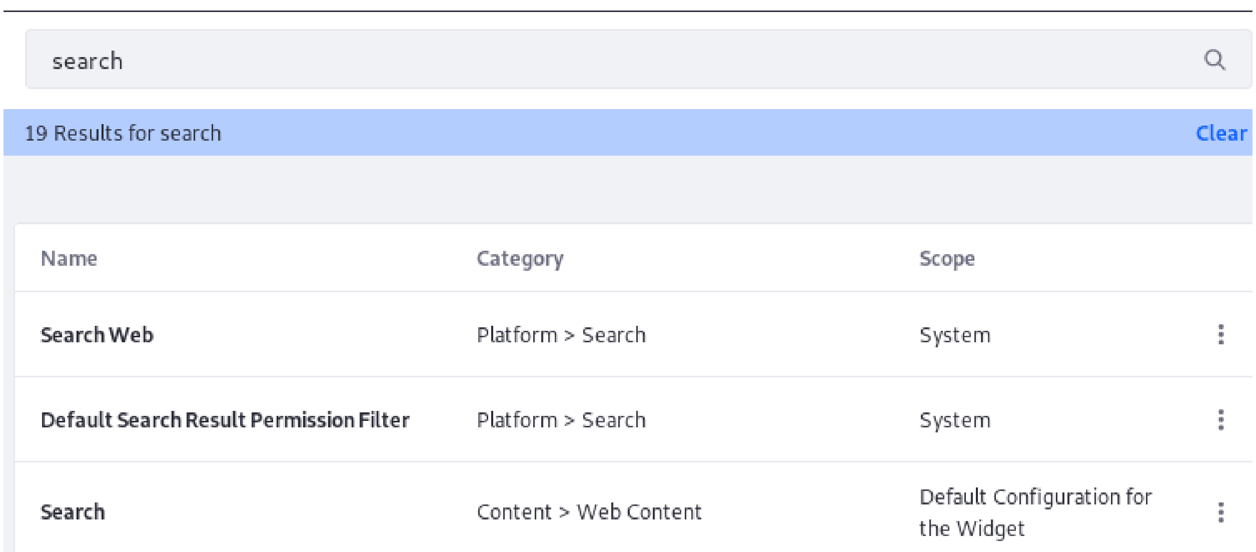

Figure 73.2: System Settings are organized by section and category.

2. Open the configuration form for the entry. Click the name of the component you want to configure.
3. Make any changes you'd like, then click *Save*. Your configuration changes are saved and applied throughout the system.

Important: Content generated using templates (e.g., FreeMarker templates and Application Display Templates) is cached. Cached content might not reflect configuration changes until the cache is invalidated (cleared). The Server Administration → Resources tab provides cache clearing options.

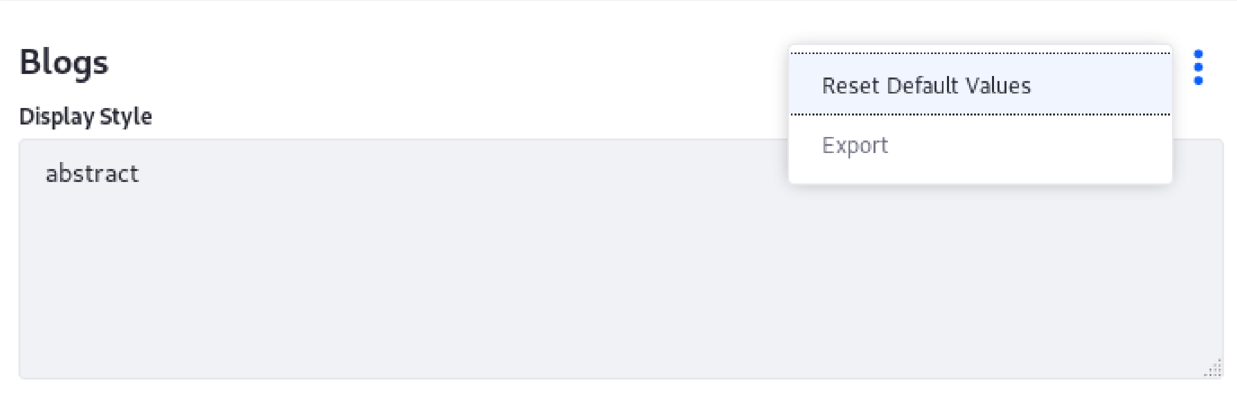

Figure 73.3: After saving changes to a configuration, the actions *Reset Default Values* and *Export* are available for it.

If you make some regrettable configuration decisions and can't recall exactly what you did, start over by clicking the actions button (⋮), then clicking *Reset Default Values*.

Configuration Scope

While browsing the categories of System Settings, you'll notice that clicking into a topic (for example, Blogs) reveals entries at different scopes. All the settings here act at the system scope. For scopes labeled other than System, these configurations act as defaults. In other words, they identify where the system-wide configuration is overridden. System scoped configurations are, of course, not overridden at more granular scopes. There are four values that you'll see under Scope:

- *System:* Any System Settings configuration made for system scoped entries becomes the final value for the application in a system-wide fashion. It affects the whole system and isn't overridden anywhere else.
- *Virtual Instance Scope:* Configuration at the Virtual Instance level is overridden in Instance Settings.
- *Site Scope:* Configuration at this scope is overridden in each site.
- *Widget Scope:* Configuration at this scope is overridden in each Widget Instance (like the Blogs example below).

Figure 73.4: Some System Settings entries are system scoped.

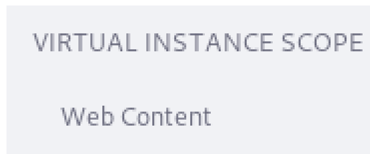

Figure 73.5: Some System Settings are virtual instance scoped.

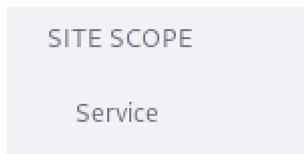

Figure 73.6: Some System Settings are site scoped.

If a configuration changed in System Settings is also configurable at a different scope, the System Settings value acts as a default that can be overridden. Once a configuration change is made at a more granular scope, making a change at the system level doesn't do anything.

For example, allowing comments is configurable for each Blogs Entry. Set the default behavior at Control Panel → Configuration → System Settings → Blogs. In the Blogs Entry under Widget Scope, disable the *Enable Comments* checkbox.

Now add a Blog Entry to a Site's Content → Blogs application. Then go to a public page and add the Blogs Widget to the page. Click the Actions button (⋮) for the widget and select *Configuration*. You'll see the same Enable Comments checkbox, and its default is now false (unchecked). Checking the box in the Widget Configuration screen breaks its link with the System Settings entry. Changing the System Settings configuration has no effect on this widget anymore.

If an application is configurable from Instance Settings and System Settings, use System Settings

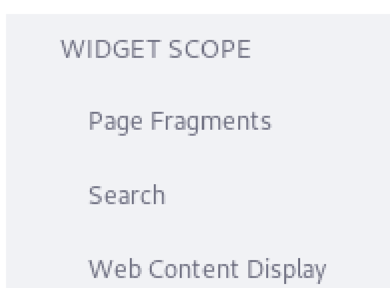

Figure 73.7: Some System Settings entries are widget scoped.

to configure it whenever possible. If you use Instance Settings and later want to revert to using the System Settings default configurations, use the *Reset Values* button from the Instance Settings entry.

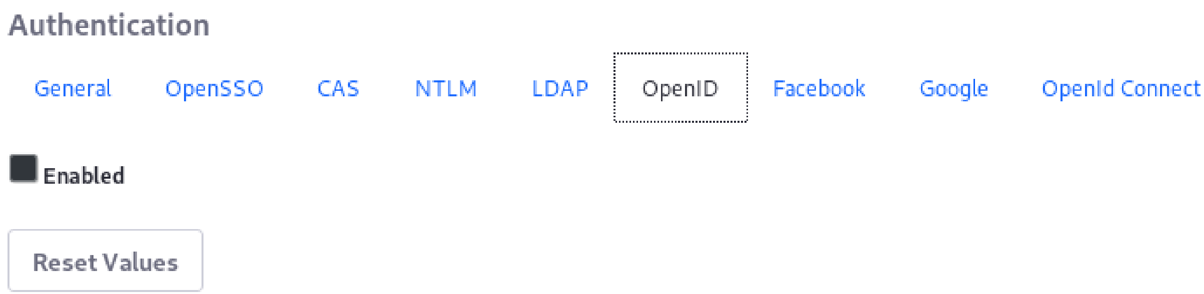

Figure 73.8: Some Instance Settings entries have a *Reset Values* button so you can safely revert your configuration changes, falling back to the System Settings defaults.

The Reset Values button removes configuration values from the database so you can rely on the defaults that were provided out-of-the-box. If there is no Reset Values button for an Instance Settings entry, once you use Instance Settings to configure the application, you can't later decide to use System Settings. Only changes to Instance Settings are recognized.

Exporting and Importing Configurations

What if you change many default configurations in System Settings, and then need to make the same changes in another installation? Don't worry, you don't need to remember and duplicate every choice you made. The System Settings application lets you export a single entry's configurations, or all the settings you made in the System Settings interface. The exported files are deployable to any other installation of the same version.

To export a single entry's configurations, click the actions button ( ), then click *Export*. A `.config` file containing your configuration downloads to your system.

To export all the configuration changes you've made in System Settings, click the System Settings options button ( ), then click *Export All Settings*. The `.config` files for all the entries you edited then download in a ZIP file.

To make these configurations active in the destination system, unzip and place the `.config` files in the `[Liferay_Home]/osgi/configs` folder.

Now you know what System Settings is and how to use it. All that's left is to explore the entries to see what configuration options are available. If you aren't sure what something does, check the documentation for the feature you're interested in, as specific configurations are covered there.

Chapter 74

Understanding System Configuration Files

The System Settings application is convenient for making system-scoped configuration changes and setting default configurations for other scopes. But there's another supported configuration approach: configuration files. You can use configuration files to transfer configurations from pre-production systems to production systems, or between any other Liferay DXP systems. Sometimes developers choose to distribute the default configuration for their applications via configuration file. Whatever the reason, configuration files offer another configuration approach.

Configuration files use the `.config` property value format defined by the Apache Felix Configuration Admin framework.

Important: Content generated using templates (e.g., FreeMarker templates and Application Display Templates) is cached. Cached content might not reflect configuration changes until the cache is invalidated (cleared). The Server Administration → Resources tab provides cache clearing options.

Note: The `.cfg` file format is common in OSGi environments, and it's a supported format, but `.config` files are preferable. You can specify a property value's type and have multi-valued properties. The syntax described in these articles is for `.config` files.

74.1 Creating Configuration Files

System Settings provides an *Export* option that becomes available once you modify a configuration entry. Exporting is the recommended way to create `.config` files: you download a `.config` file containing the entry's settings in a `key=value` format. Liferay DXP exports an entry's total available configuration keys and values, even if only one value was changed. You can export a single configuration entry or the entire set of modified configurations.

To avoid a file name conflict, name configuration files using a unique identifier. For example, the Journal Service entry, which backs Web Content functionality, has this file name:

```
com.liferay.journal.configuration.JournalServiceConfiguration.config
```

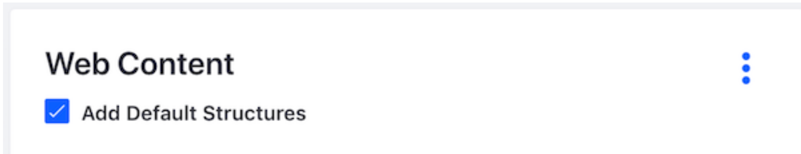

Figure 74.1: The Web Content System Settings entry has the back-end ID `com.liferay.journal.configuration.JournalServiceConfiguration`.

Key/Value Syntax

The syntax for all keys and values in a `.config` file is the same:

```
configurationName="value"
```

For single value configurations without special characters, that's all there is to know. Settings with multiple values and certain characters require slight modifications.

Multi-Value Settings

Configuration entries can have properties that accept multiple values. For example, a configuration property for specifying supported file extensions needs more than one value. Here's how to write a multi-value setting in a `.config` file:

```
multiValueSetting=["Value 1","Value 2", ...]
```

Do not use a space character between values (after the comma). The property won't be loaded.

Open the Web Content category in System Settings (under the Content section), and select *Web Content* for the virtual instance scope. You'll see what looks like multiple single value entries for *Characters Blacklist*:

In the configuration file, this is really a single key with an array of comma-separated values:

```
charactersblacklist=["&","'","@","\\","]","}",":","\=",">","/","<","[","{","%","+","#","`","?","\"",";","*","~"]
```

Escaping Characters

Double quotes (") and equals signs (=) must be *escaped* in `.config` files. Escaping is using another character to denote that a character shouldn't be used in its normal way. Since double quotes and equals signs are already used in `.config` files, escaping them tells the framework not to read them the normal way, but to pass them through as part of the value. Use a \\ to escape characters in the `.config` file:

```
charactersblacklist=["&","\"","\="]
```

This setting illustrates a multi-value setting with a regular, unescaped character (&), and two escaped ones (\" and \=).

Along with the mandatory escaping of double quotes and equals characters, it's beneficial to escape spaces inside values to avoid problems. In this example, a \\ is used before each space character to ensure it's read and processed properly:

```
blacklistBundleSymbolicNames=["Liferay\ Marketplace","Liferay\ Sharepoint\ Connector"]
```

If you don't escape spaces yourself, the framework adds the backslash for you after deployment.

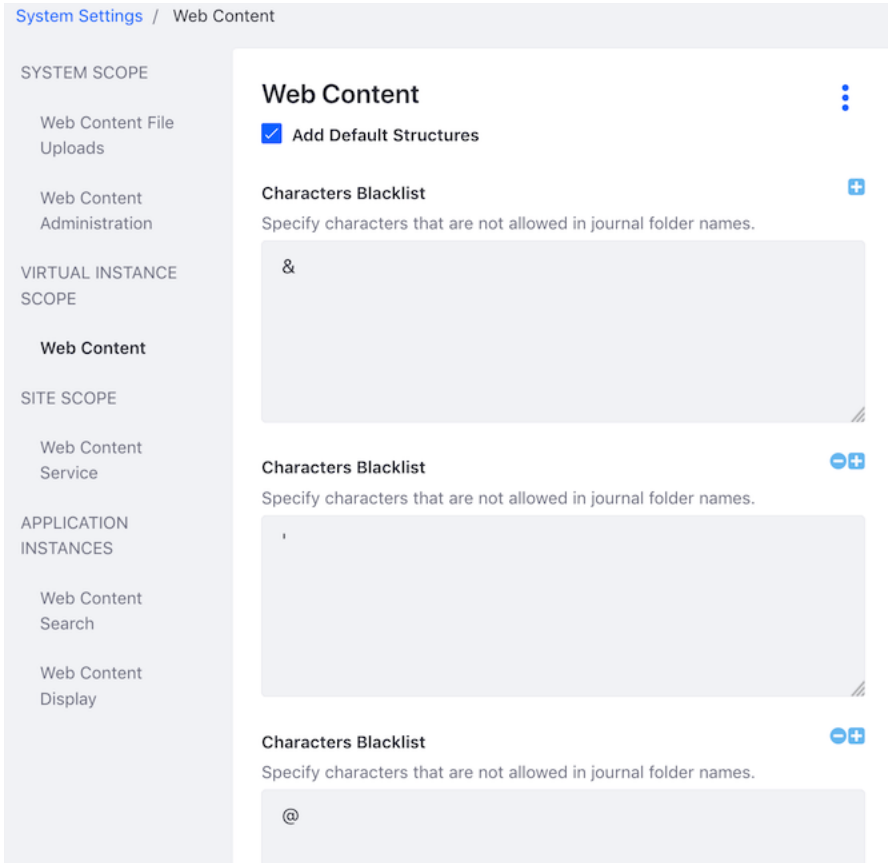

Figure 74.2: The Web Content System Settings entry has many *Characters Blacklist* fields.

Typed Values

The `.config` file format supports specifying the type of a configuration value by inserting a special type marker character. Because Liferay DXP already knows the correct type for each configuration property, the type characters are only useful for informational purposes. For example, a configuration with a boolean type has *B* just before the value to mark it as a boolean type:

```
addDefaultStructures=B"true"
```

If you see type markers in `.config` files, you can safely ignore them. The example included above functions identically without the type marker:

```
addDefaultStructures="true"
```

Deploying a Configuration File

Once you have a configuration file, deploy it. It's registered and the targeted configuration values are updated automatically.

To deploy the `.config` file, place it in your Liferay Home's `osgi/configs` folder. To change the configuration further, you can edit the `.config` file directly or use System Settings.

Configuration Files and Clustering

In a clustered environment, each node needs the same system configuration. If you're assembling an image of Liferay on which to base cluster nodes, you can export your System Settings to `.config` files to your image's `[Liferay Home]/osgi/configs` folder. The configurations are applied to each cluster node you create.

74.2 Factory Configurations

Configurations supporting multiple entries are called *factory configurations*.

Factory Configuration Example: JAX-WS and JAX-RS web services are supported. These services must use a CXF Endpoint, which is a context path where the web services are deployed and accessed. Endpoints can be created via factory configuration by navigating to the CXF Endpoints entry in System Settings (System Settings → Platform → Web API → CXF Endpoints). Click *ADD*, enter the desired configuration values, then repeat the process to add as many CXF Endpoint configurations as needed. Creating CXF Endpoint configurations also creates CXF Endpoints themselves. This is how factory configurations work.

If a service is meant to support factory configurations, its System Settings entry has an ADD button.

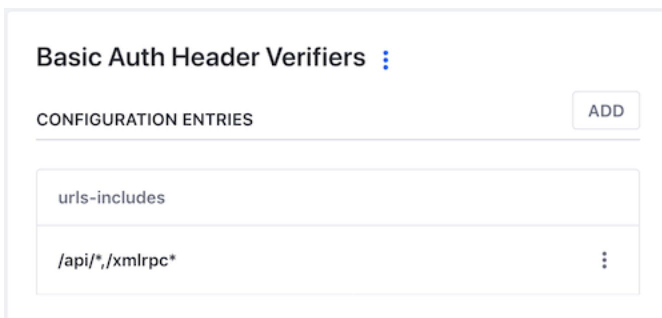

Figure 74.3: If a System Settings entry has an ADD button, it's suitable for factory configurations.

As with single-instance configurations, you can set factory configurations in the System Settings interface (as described in the example above) or via configuration files. Name a standard single-instance configuration file like this:

```
my.service.ServiceConfiguration.config
```

If your service supports factory configurations, use the convention of calling the configuration's first instance `-default.config`:

```
my.service.ServiceConfiguration-default.config
```

The next instance contains a unique *subname* (something other than *default*). It's good practice to use a descriptive name:

```
my.service.ServiceConfiguration-port9080.config
```

To follow the CXF Endpoints example described above, if Liferay's developers had shipped an initial CXF Endpoint `.config` file with Liferay DXP, it would have been named this:

```
com.liferay.portal.remote.cxf.common.configuration.CXFEndpointPublisherConfiguration-default.config
```

If this `-default.config` configuration specifies a context path for *REST* web services, and you create another endpoint with a different context path for *SOAP* web services, your second configuration file could be named:

```
com.liferay.portal.remote.cxf.common.configuration.CXFEndpointPublisherConfiguration-soap.config
```

Note: Some System Settings entries (like the CXF Endpoints entry) don't ship with a configuration file, so anything you create is the first occurrence. However, if you configure one and export it to obtain the `.config` file, it doesn't use the `-default.config` naming convention. Instead it's given a guaranteed unique identifier for its subname, like this:

```
com.liferay.portal.remote.cxf.common.configuration.CXFEndpointPublisherConfiguration-a6f67e48-6dca-49c6-bf6b-8fd5e6016b2d.config
```

This guarantees that the file has a unique name. If you're exporting the configuration file for deployment in a separate system, you can rename the exported file to use a more descriptive subname. If you rename the file and deploy it to the same system it was exported from, the new subname marks it as an entirely new configuration. You'll end up with an additional configuration instance in this case, not just a renamed one.

Warning:: For configuration entries supporting factory configurations, omitting the subname from a `.config` file's name causes System Settings to disallow adding new entries for the configuration (only the configuration entry targeted by this `.config` file). This is caused by a known bug. See LPS-76352 for more information. Once an improperly named configuration file is deployed, you can't add any entries for the configuration in question from its System Settings entry. For example, if you deploy the following file to configure a CXF Endpoint, not only does this not add a CXF Endpoint, it also prevents you from adding any CXF Endpoints via System Settings:

```
com.liferay.portal.remote.cxf.common.configuration.CXFEndpointPublisherConfiguration.config
```

Deploying an erroneous (lacking a subname) `.config` file doesn't disable anything permanently. Just rename the file using the proper convention described above or remove it entirely and start over.

In many cases, configuration files can be used to force a factory configuration scenario, but not all configurations can be used this way. It's best to stick to the intended use cases. Use System Settings as described above to determine if using factory configurations is a good idea. If not, stick to the single occurrence mode of configuration (specifying only one configuration file for the service).

Chapter 75

Server Administration

Server Administration lets you manage and monitor your Liferay DXP server. Access the application by clicking *Control Panel → Configuration → Server Administration*.

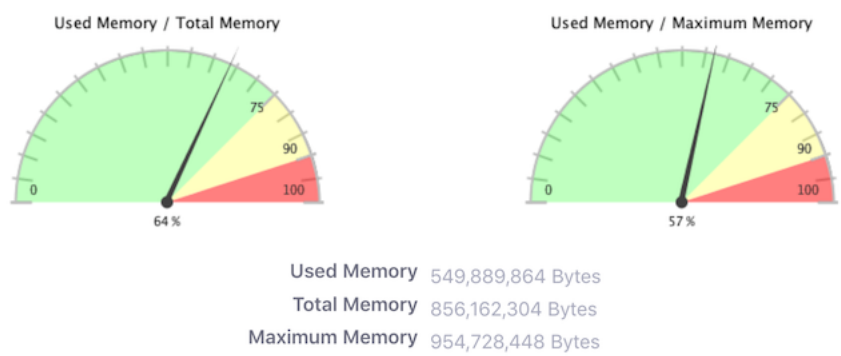

Figure 75.1: The Resources tab of Server Administration shows a graph of your server's memory usage.

Server Administration's functionality is segmented into these tabs:

Resources: View memory usage and perform management tasks like running the garbage collector, clearing the database cache, and more. For more information, see Resources.

Log Levels: View and set logging levels. You can make dynamic modifications of log levels for any class hierarchy in Liferay DXP. Custom objects not on the list can be added with the *Add Category* tab. Changes to the log level near the top of the class hierarchy (such as at `com.liferay`) also change log levels for all the classes under that hierarchy. Modifications unnecessarily high in the hierarchy generate too many messages to be useful.

Properties: View JVM and portal properties. This tab has two sub-tabs: System Properties and Portal Properties. The System Properties tab shows an exhaustive list of system properties for the JVM, as well as many Liferay DXP system properties. You can use this information for debugging purposes or to check the currently running configuration. The Portal Properties tab shows an exhaustive list of the current portal property values. For explanations of these properties, see the portal properties reference documentation.

Data Migration: Migrate documents from one repository to another. For example, you can migrate your documents to a new repository on a different disk or in a new format. Here are the steps:

1. Create a backup copy of the Document Library repository and Liferay DXP database.
2. Configure the new file store in *System Settings → Platform: File Storage*.
3. In this tab (*Server Administration → Data Migration*), select the repository hook for the file store you configured and click *Execute*.
4. Make sure the data migrated correctly.
5. Configure the new repository as the default.
6. If you used a `portal-ext.properties` file to configure the repository, restart the server.

Mail: Instead of using a `portal-ext.properties` file to configure a mail server, you can configure a mail server from this tab. If your message boards receive mail, you can connect a POP mail server. If Liferay DXP sends mail (useful for sending notifications to users), you can connect to an SMTP server. Note that if you configure mail server settings here in System Settings, these settings override any mail server settings in your `portal-ext.properties` file.

External Services: Configure external services for generating file previews. For more information, see the article on External Services.

Script: A scripting console for executing migration or management code. The Groovy scripting language is supported out of the box. .

Shutdown: Schedule a shutdown that notifies logged-in users of the impending shutdown. You can define the number of minutes until the shutdown and a message to display. Liferay DXP displays the message at the top of users' pages for the duration of time you specified. When the time expires, all pages display a message saying the portal has been shut down. The server must then be restarted to restore access.

75.1 Server Administration: Resources

The Server Administration app's Resources tab contains several server wide actions that an administrator can execute. These include the following items:

Run the garbage collector: Tells the JVM to free memory by running the garbage collector.

Generate a thread dump: Generates a thread dump for later scrutiny to determine the presence and location of any deadlocks. Useful during performance testing, but you must add a logger category for `com.liferay.server.admin.web.internal.portlet.action.EditServerMVCActionCommand` and set it to `INFO` before executing.

Clear content cached by this VM: Clears content stored in the local cache. Only local JVM scope Ehcache content is cleared, not clustered Ehcache. 1

Clear content cached across the cluster: Clears the content of the entire clustered Ehcache. 1

Clear the database cache: Clears the database cache. Does not clear any Ehcache content except database results at the persistence layer. 1

Clear the direct servlet cache: Clears the direct servlet cache. In case emergency fixes must be applied, this action allows an administrator to clear out the cache manually to force JSPs to reload.

The direct servlet context optimizes JSP serving performance by caching and accessing the generated servlets directly instead of accessing them over the application server's dispatcher chain. This function is only suitable for cases where no filter is required for the JSPs; it should be enabled for production mode to improve performance, but disabled for development mode to allow JSP servlets to be reloaded on the fly. See the Direct Servlet Context section of the `portal.properties` file for details. 1

Verify database tables of all plugins: Checks all tables against their indexes for data retrieval accuracy.

Verify membership policies: Checks that existing Site membership policies were correctly applied and automatically makes updates. If the Liferay DXP database is changed manually or is hacked—resulting in a user assigned to a Site in violation of a site membership policy—this action triggers the verification methods of all implemented Site membership policies. Changes are automatically made to any memberships in violation.

Reset preview and thumbnail files for Documents and Media: Regenerates previews of each item in your Documents and Media libraries.

Clean up permissions: Removes permissions on the Guest, User, and Power User Roles to simplify the management of User Customizable Pages. The Add To Page permission is removed from the Guest and User Roles for all portlets, and is reduced in scope for Power Users from portal-wide to User Personal Site.

Clean up portlet preferences: This action cleans up database entries if portlet preferences become orphaned in the Liferay DXP database.

[1] Caching occurs at multiple levels. Some higher caching layers aren't aware of lower caching layers. Always clear the cache at the lowest (most granular) layer possible, even if you've already cleared a higher level cache.

75.2 Server Administration: External Services

Users can upload and share any type of file via the Documents and Media library, a customizable and permissions-enabled online repository for files (see publishing files for more information). PDFBox is included with Liferay DXP and generates automatic previews for certain file types (mostly PDFs). You can install two additional tools to generate previews for other file types:

- **OpenOffice:** or **LibreOffice:** Convert and generate previews for many types of documents.
- **ImageMagick:** Generate higher-quality image previews for many types of images.
- **Xuggler:** Convert and generate previews for audio and video files.

As of Liferay 7.1, OpenOffice/LibreOffice is configured via OSGi Configuration Admin instead of via portal properties. Thus, to adjust the these settings, you have to go to *Control Panel* &arr; *Configuration*&arr; *System Settings* &arr; *Other* → *OpenOffice Integration*. You can also adjust these settings via a `com.liferay.document.library.document.conversion.internal.configuration.OpenOfficeConfiguration.config` file in your `[Liferay Home]/osgi/configs` folder. See the Breaking Changes document for details.

Once you've installed ImageMagick and Xuggler, you can use the Server Administration app's External Services tab to configure Liferay DXP to use them. Make sure to choose the correct versions of these tools for your operating system. We recommend that you install the latest stable versions, as older versions may not run properly with Liferay DXP. ImageMagick must be installed manually, but you can install Xuggler from the Control Panel.

Tip: If you're running Liferay DXP on a Linux server and experience a problem enabling Xuggler, check your server's glibc version. For Xuggler to work, you may need to update glibc to version 2.6 or later.

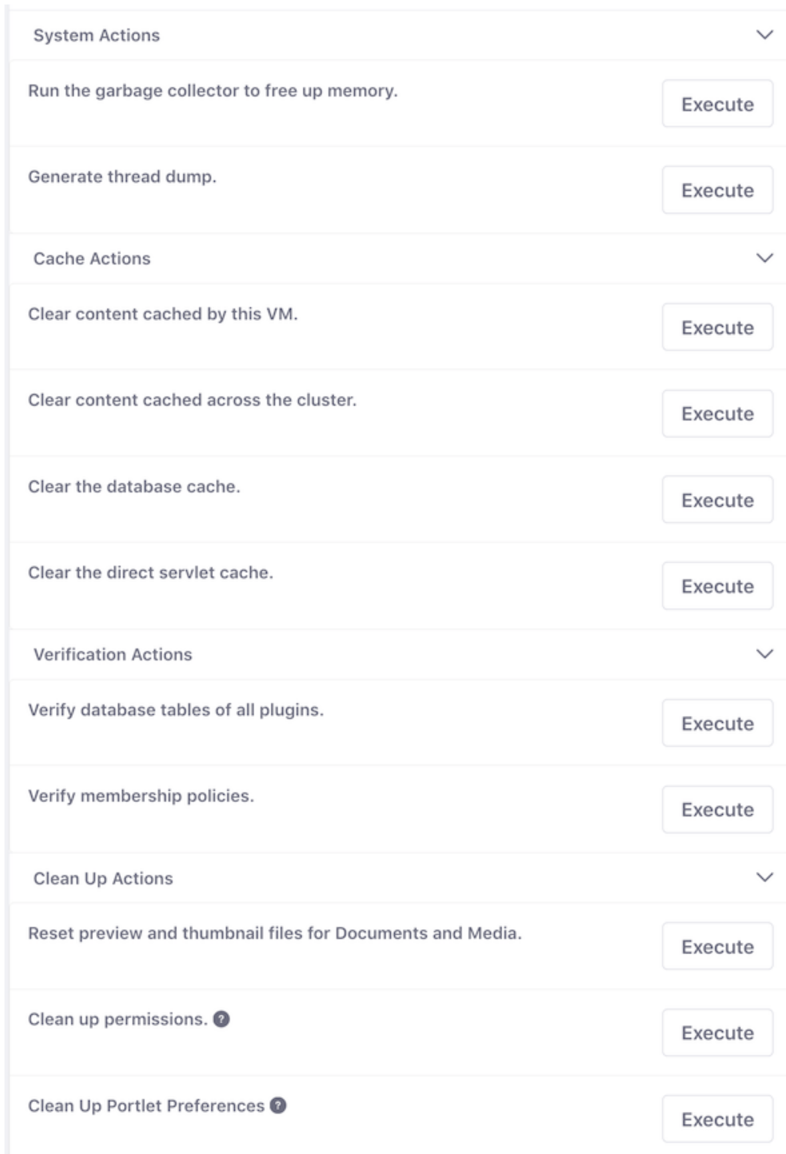

Figure 75.2: The Resources tab of Server Administration lets you execute several server maintenance tasks.

ImageMagick Configuration

Before configuring ImageMagick to generate image and PDF previews, install it and its dependency, Ghostscript. This differs by operating system: on Linux, both are likely already installed. They are not likely to be installed on Windows, but may be on macOS.

1. Download and install ImageMagick.
2. Download and install Ghostscript.

Once installed, you must enable ImageMagick in the Server Administration app's External Services tab, or in a `portal-ext.properties` file. If using `portal-ext.properties`, add the following lines and make sure the search path points to the directories containing the ImageMagick and Ghostscript executables. You may also need to configure the path for fonts used by Ghostscript when in macOS or Unix environments.

```
imagemagick.enabled=true
imagemagick.global.search.path[apple]=/opt/local/bin:/opt/local/share/ghostscript/fonts:/opt/local/share/fonts/urw-fonts
imagemagick.global.search.path[unix]=/usr/local/bin:/usr/local/share/ghostscript/fonts:/usr/local/share/fonts/urw-fonts
imagemagick.global.search.path[windows]=C:\\Program Files\\ImageMagick
```

Follow these steps to instead enable ImageMagick from the Server Administration app's External Services tab:

1. In the *Control Panel*, navigate to *Configuration* → *Server Administration*, then click the *External Services* tab.
2. Expand the ImageMagick and Ghostscript section and select *Enabled*.
3. Verify that the paths to the ImageMagick and Ghostscript executables are correct.

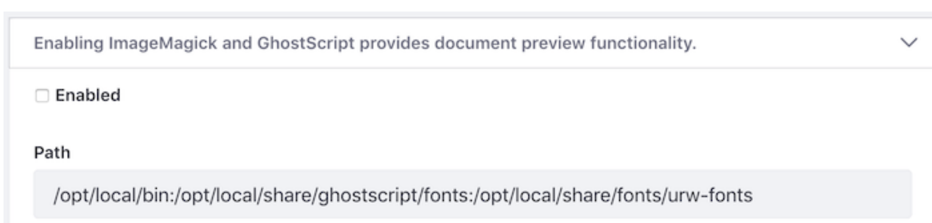

Figure 75.3: Enable ImageMagick and Ghostscript, and verify that the paths are correct.

Xuggler Configuration

Follow these steps to install and configure Xuggler:

1. In the *Control Panel*, navigate to *Configuration* → *Server Administration*, then click the *External Services* tab.
2. In the Xuggler section, select the Xuggler `.jar` file that matches your operating system. Then click *Install*.
3. Restart your application server.

4. Enable Xuggler for your portal. There are two ways to do this:

    - In the Control Panel, navigate to the *Server Administration* → *External Services* tab, select *Enabled*, then click *Save*.
    - Add the following line to your `portal-ext.properties` file and restart your application server:

```
xuggler.enabled=true
```

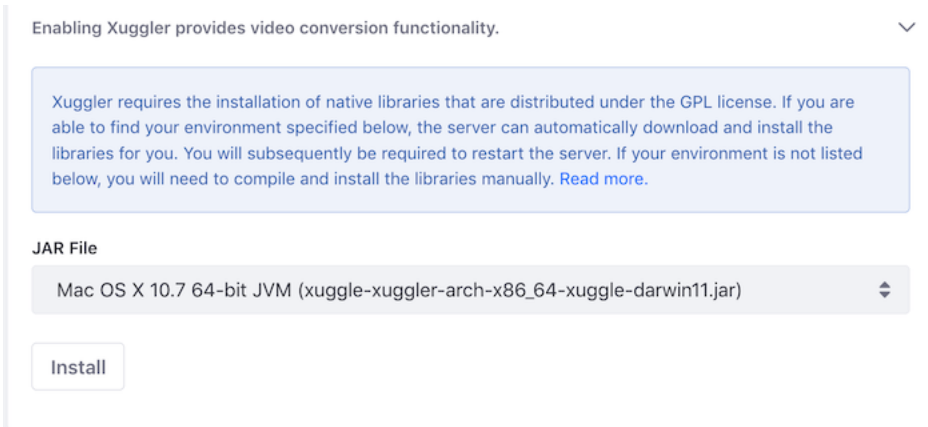

Figure 75.4: Install Xuggler.

Chapter 76

Setting Up a Virtual Instance

Once Liferay DXP is installed, the configuration begins. Recall that configuration happens at different scopes. Here we're covering configuration at the virtual instance scope. There's an important difference between the system scope and the instance scope. The system scope is the highest level scope you can make configurations at. All virtual instances are impacted by configuration done at this scope. The instance scope applies only to one particular virtual instance.

Virtual instances have unique domain names but share a server and database. Each virtual instance can have independent data and configurations.

These articles cover

- Adding a virtual instance
- Configuring a virtual instance

Get started by learning about adding a virtual instance.

76.1 Virtual Instances

Here's a quick scenario: you already have a server hosting a Liferay DXP installation and a database. It has many Users, Sites, and specific instance settings. If you need another such installation, then adding a *Virtual Instance* might be right for you.

You can run more than one Virtual Instance on a single server with a shared database, but separate data and configurations. Users are directed to the correct Virtual Instance via its unique domain name. And because Virtual Instances share an application server and OSGi container, they also share these customizations:

- Custom code deployed by developers and administrators.
- System-scoped configurations (e.g., `.config` files, changes made in *Control Panel → Configuration → System Settings*).
- Application server configuration.

Administrators can manage Virtual Instances in *Control Panel → Configuration → System Settings*. Follow these steps to create a Virtual Instance:

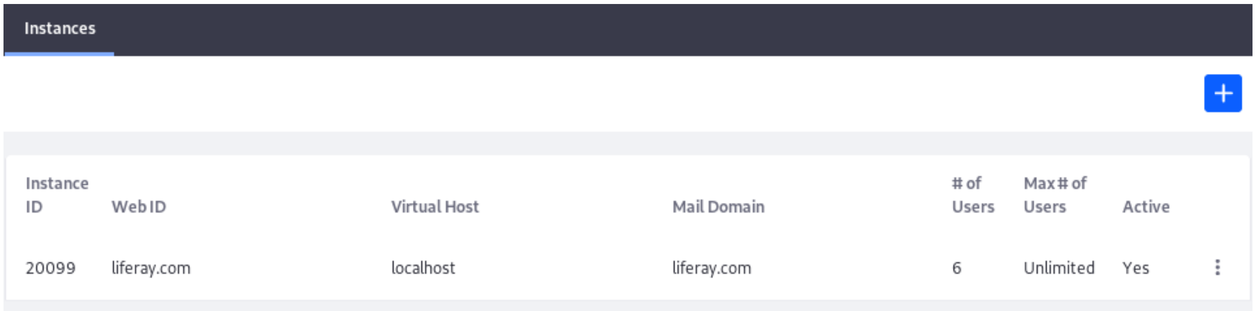

Figure 76.1: Add and manage virtual instances of Liferay in the Control Panel's *Configuration* → *Virtual Instances* section.

1. Before you create a Virtual Instance, configure its domain name in your network.

2. Go to *Control Panel* → *Configuration* → *Virtual Instances*.

3. Click the *Add* button (+). This opens the *New Instance* form.

4. Complete the New Instance form as follows:

Web ID: The instance's ID. Using the domain name is a common convention.

Virtual Host: The domain name you configured in your network. When users are directed to your server via this domain name, they'll be sent to the Virtual Instance that contains their data.

Mail Domain: The mail host's domain name for the Virtual Instance. Email notifications are sent from the instance using this domain.

Max Users: The maximum number of user accounts the Virtual Instance supports. Enter *0* to support unlimited users.

Active: Whether the Virtual Instance is active. Note that inactive Virtual Instances aren't accessible to anyone, even the administrator.

Click *Save* when you're done filling out the form.

Now you can navigate to the instance using its domain name. You're brought to what looks like a clean install of Liferay DXP. This is your new Virtual Instance! You can configure it any way you like. The remaining articles in this section show you how to configure an instance's settings.

76.2 Terms of Use

Terms of Use are important when you need them, but not all Sites do. They're not first in the Instance Settings interface, but they're one of the first things to configure for your virtual instance.

Configure Terms of Use for the virtual instance at *Control Panel* → *Configuration* → *Instance Settings*. In the Configuration tab, scroll down to the Terms of Use section and expand it.

Since the Terms of Use page is enabled by default, one of your first actions should be to disable or replace the default, placeholder terms. You can disable the requirement for all Users to read the terms of use or set the Group ID and Article ID for the Web Content Article that contains your Terms of Use.

To find a Web Content Article's ID:

1. Go to the *Site Menu* → *Content* → *Web Content*.

2. Click on the Terms of Use article. Its ID appears at the top of the screen, with the Version and Workflow Status.

Figure 76.2: The Web Content Article ID is displayed in the edit screen.

The Group ID is the ID of the Site the Web Content is associated with. To find the Group/Site ID:

1. Go to *Site Menu* → *Configuration* → *Site Settings*.
2. Find the Site ID field in the General tab. Enter it into the Group ID field.

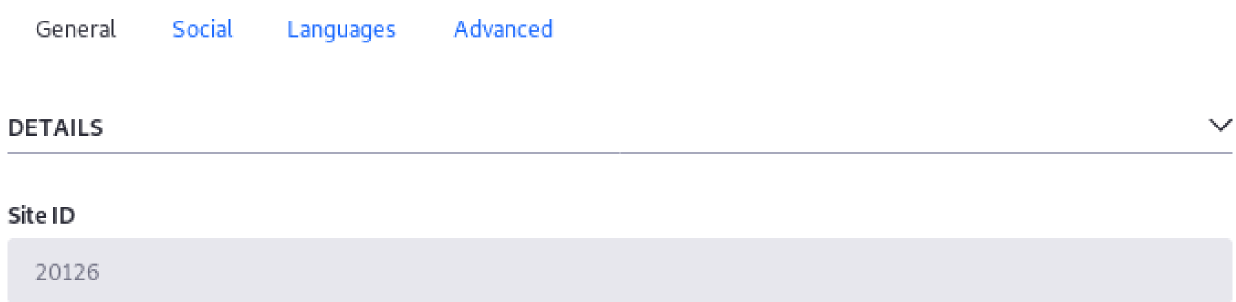

Figure 76.3: The Site ID in Site Settings is the Group ID in the terms of Use configuration.

Save the configuration and now all users signing in for the first time see your Terms of Use article. Users must agree to the Terms of Use to register their User Accounts.

76.3 General Settings

General virtual instance settings are configured at *Control Panel* → *Configuration* → *Instance Settings*. The General settings are at the top of the first tab, Configuration.

The General section has three headings: Main Configuration, Navigation, and Additional Information. Under Main Configuration,

- Set the name of the entity responsible for running the virtual instance.
- Set the mail domain.
- Set the virtual host.
- Set the URLs to a CDN configured to serve static resources.

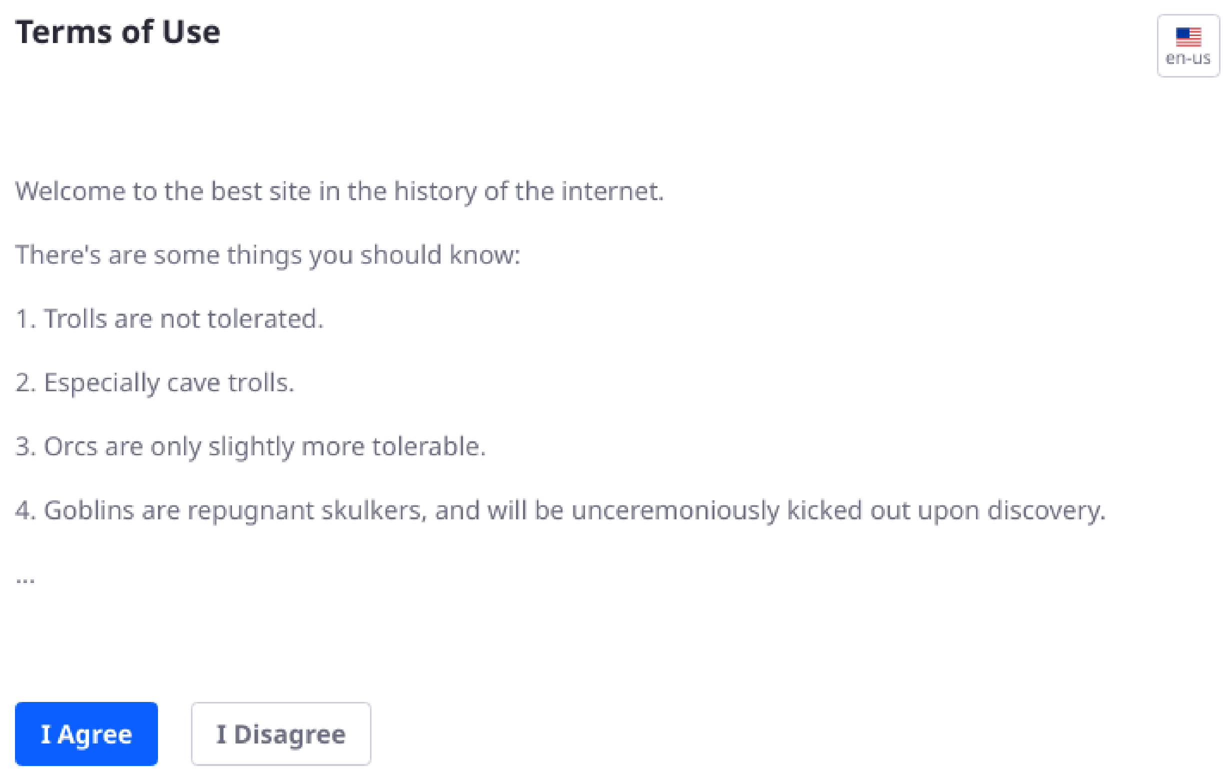

Figure 76.4: Turn a Web Content Article into the Site's Terms of Use page.

Under the Navigation heading, set a home page for your virtual instance as well as default landing and logout pages. To set these pages, use the part of the page's address that follows your domain. For example, to set the default landing page to be `http://localhost:8080/web/guest/login`, use `/web/guest/login`. Use the variable `${liferay:screenName}` as part of the address. This comes in handy if you want to redirect Users to their personal pages upon login. For example, the User `marvin` has this default URL to his personal page:

```
http://localhost:8080/user/marvin
```

To make sure he's directed there on login, place this in the Default Landing Page field:

```
/user/${liferay:screenName}
```

These URLs can also be set at the system scope, in a `portal-ext.properties` file:

```
default.landing.page.path=
default.logout.page.path=
company.default.home.url=
```

All virtual instances share the values specified in the properties file. Changes made in Instance Settings override the values set in the properties file.

For more information, see the Portal Properties documentation.

Under the Additional Information heading, you can specify a Legal name, ID, company type, SIC code, ticker symbol, industry and industry type.

76.4 Authentication

Configure authentication at Control Panel → Configuration → Instance Settings. In the Configuration tab, scroll down to the Authentication section and expand it.

Authentication has several tabs: General, OpenSSO, CAS, NTLM, LDAP, OpenID, Facebook, Google, and OpenId Connect. These configure how Users authenticate.

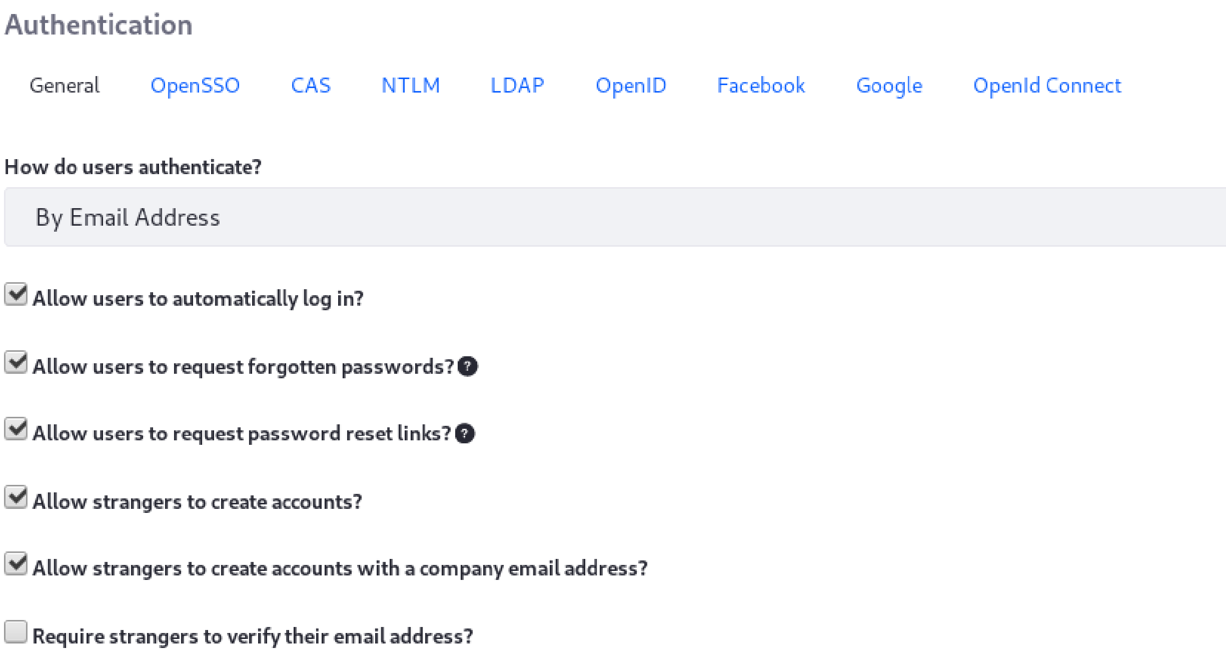

Figure 76.5: Configure general authentication behavior and settings for external authentication systems.

The General tab's settings affect only the platform's built-in authentication functionality and don't have anything to do with the integration options on the other tabs. There are several general authentication settings:

- Authenticate by email address (default), screen name, or User ID (a numerical ID auto-generated in the database—not recommended).
- Enable/Disable automatic log in. If enabled, a User can check a box which will cause the Site to "remember" the login information by placing a cookie on the browser. If disabled, Users must always log in manually.
- Enable/Disable forgotten password functionality.
- Enable/Disable request password reset links.
- Enable/Disable account creation by strangers. If running an Internet site, leave this enabled so visitors can create accounts on your Site.
- Enable/Disable account creation by those using an email address in the domain of the company running the Site (which is set on the General page of Instance Settings). This is handy if

you're using Liferay to host both internal and external web sites. Make sure all internal IDs are created by administrators but external Users can register for IDs themselves.

- Enable/Disable email address verification. If enabled, Users receive a verification email with a link back to the virtual instance, verifying that the email address they entered is valid.

By default, all settings except for the last are enabled. User authentication by email address is an important default for the following reasons:

1. An email address is unique to the User who owns it.
2. People remember their email addresses. A Users who hasn't logged in for a while could forget their screen name.
3. If email address isn't used to authenticate, a User might change her email address then forget to update the email address in her profile. If this occurs, no notifications sent by the virtual instance will reach the User. Keeping the email address at the forefront of a User's mind when she logs in helps ensure the User keeps it current.

Learn to configure a third party authentication service or set up Single Sign On (SSO) in the security documentation.

76.5 Users

Configure user settings at Control Panel → Configuration → Instance Settings. In the Configuration tab, scroll down to the Users section and expand it.

The *Users* section has three tabs: Fields, Reserved Credentials and Default User Associations.

In the Fields tab, enable/disable these fields on the Add/Edit User Form:

- Autogeneration of screen names
- Birthday field
- Gender field

In the Reserved Credentials tab, specify the screen names and email addresses Users aren't allowed to use. This prevents Users from creating IDs that look like administrative IDs or that have reserved words in their names.

Default User Associations has four fields for listing (one per line) the Sites, Organization Sites, Roles, and User Groups you want all new Users assigned to automatically. By default, new Users are assigned to only the Users Role. You can also choose whether to apply these configurations to existing Users.

User groups are handy tools for pre-populating your Users' private Sites with pages, assigning Roles and permissions, and managing site membership.

76.6 Email and Content

There are a few sections of the Configuration tab dedicated to configuring email and content sharing:

- Mail Host Names
- Email notifications
- Content Sharing

Configure these settings at Control Panel → Configuration → Instance Settings. In the Configuration tab, scroll down to the bottom. These three entries are last in the tab.

Mail Host Names

Enter one mail host name per line, besides the one you configured on the General tab. This tells the virtual instance which mail host names are owned by your organization.

Email Notifications

The Sender tab sets the virtual instance's administrative name and email address. By default, they are from the `admin.email.from.name` and `admin.email.from.address` portal properties. This name and email address appear in the *From* field in all email messages sent by the virtual instance.

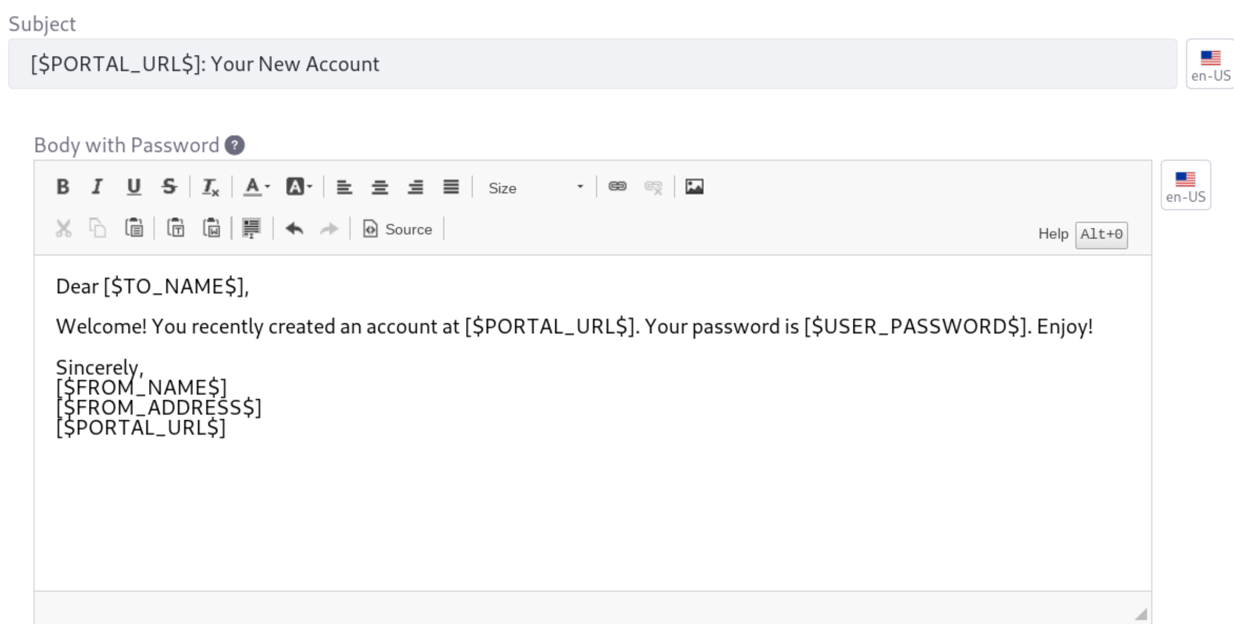

Figure 76.6: Customize the email template for the email messages sent to new Users.

The remaining tabs (Account Created Notification, Email Verification Notification, Password Changed Notification and Password Reset Notification) customize the email messages sent to Users each time any of those four events occur.

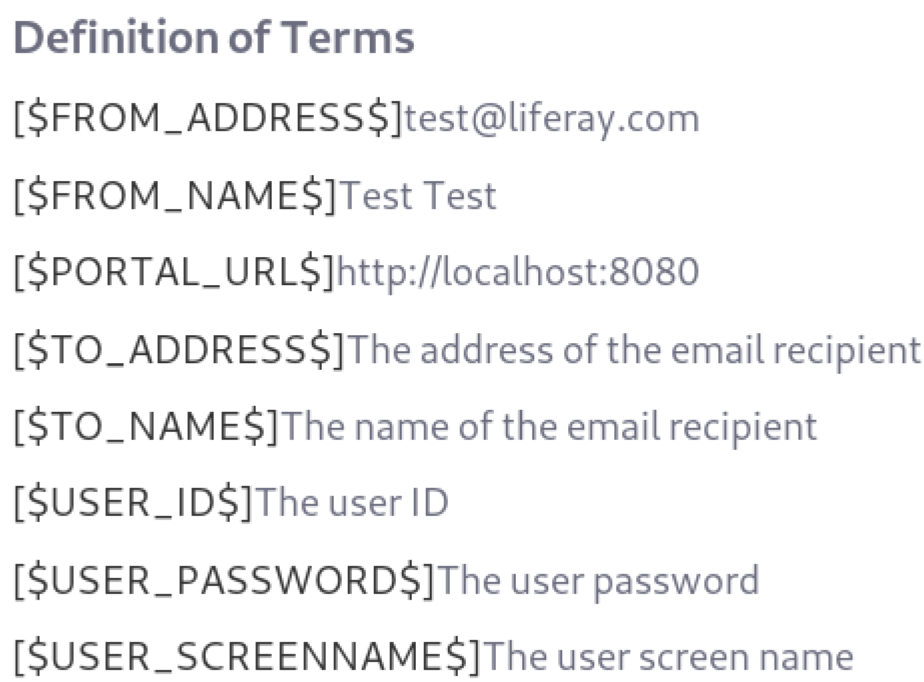

Figure 76.7: There are some handy variables available for use in email templates.

A list of tokens, entitled "Definition of Terms," is provided for insertion of certain variables (such as the portal URL or the User ID) into custom email messages.

Content Sharing

Choose if Site administrators can display content in Sites from other Sites they administer. For example, suppose that a certain User is a Site administrator of two Sites: *Engineering* and *Marketing*. The checkbox in the Content Sharing section of Instance Settings determines if the Site administrator can display content from the Marketing Site in the Engineering Site and vice versa.

You can also choose if sub-sites can display content from parent Sites and configure the defaults. There are three options:

Enabled by Default: Subsites can display content from parent Sites by default, but this can be disabled by a Site administrator.

Disabled by Default: Subsites cannot display content from parent Sites by default, but this can be enabled by a Site administrator.

Disabled: Subsites cannot display content from parent Sites, and this behavior cannot be changed by a Site administrator.

That covers a lot of Instance Settings, but you're not finished yet. The next article covers the identification and social settings.

76.7 Identification and Social

The Identification and Social sections of Instance Settings are smaller than some of the other sections, but contain important settings for some virtual instances. Both are top level tabs of

Control Panel → Configuration → Instance Settings.

Identification

The Identification section has several links for addresses, phone numbers and other contact information for the organization that owns the instance. Developers can query for this information in their applications.

Social

The Social section configures Ratings and Mentions throughout the instance. For each portlet with ratings enabled, choose whether to use a simple *Like*, *Thumbs*, or *Stars* for rating the content. You can also choose whether Users can mention (*@Username*) other Users, and if so, whether all Users can do that or you want to further refine the mentions configuration.

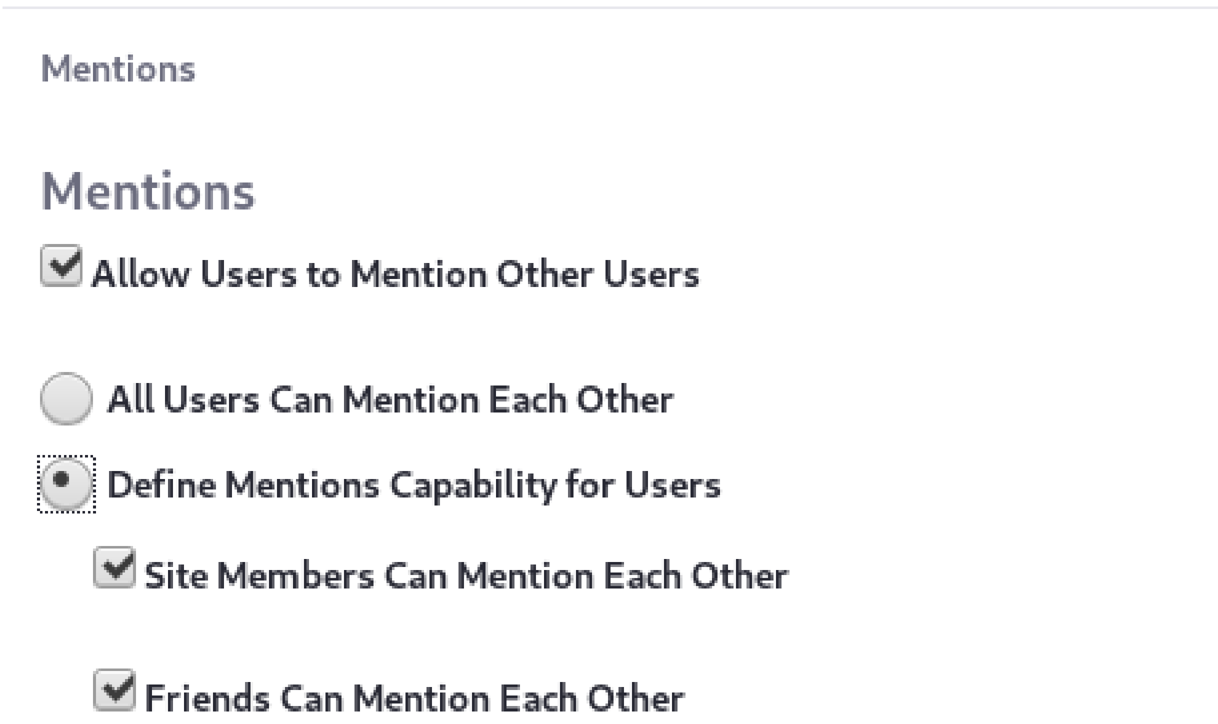

Figure 76.8: Configure the Mentions capability throughout a virtual instance.

76.8 Miscellaneous Settings

The Miscellaneous tab of Control Panel → Configuration → Instance Settings contains display settings and a few integration configurations.

Miscellaneous: Display Settings

Set the default instance language and time zone. You can also choose an instance-wide logo that appears in the top left corners of pages.

When configuring a new logo, be careful to choose an image file that fits the space. Large images might overlap with the navigation.

Miscellaneous: Integrations

In the Miscellaneous tab below the Display Settings form, there are three additional configuration forms for these integrations:

Analytics: Enter an analytics system or remove one of the two pre-configured options (`google` and `piwik`). Activate these systems here, and configure them at the site level.

Maps: Choose OpenStreetMap or Google Maps as the maps API provider for geolocalized asset display.

Google Apps: Specify the Google Apps API Key and Google Client ID to integrate Google Apps.

Chapter 77

Using Liferay's Script Engine

Liferay DXP provides a robust script engine for executing Groovy scripts. You can execute scripts to perform maintenance tasks like data cleanup, user maintenance operations, bulk invocations of Liferay's API, or even system level operations in the scripting console.

These tutorials cover the following scripting topics:

- Invoking Liferay services
- Running scripts from the script console
- Using the script engine with workflow
- Script examples

The most common use of the scripting console is to invoke Liferay's services, so that's covered first. Familiarity with Liferay's API makes the scripting console a useful tool.

77.1 Invoking Liferay Services From Scripts

Many scripting scenarios require invoking Liferay services. Liferay `*ServiceUtil` classes are the fastest and most convenient way to invoke Liferay services in the script console. You can use Groovy to invoke Liferay services the same way you would use Java. Groovy's syntax facilitates writing concise elegant scripts.

This first example illustrates correct syntax for interacting with Liferay services. It uses `UserLocalServiceUtil` to retrieve a list of users and print their names to Liferay's log. To do this, you could deploy a module with Java code like this:

```
import com.liferay.portal.kernel.model.User;
import com.liferay.portal.kernel.service.UserLocalServiceUtil;
import java.util.List;

...

int userCount = UserLocalServiceUtil.getUsersCount();
List<User> users = UserLocalServiceUtil.getUsers(0, userCount);
```

```
for (User user:users) {
    System.out.println("User Name: " + user.getFullName());
}
```

...

Or you could use Groovy—based on Java—and do the whole thing right from the script console with the same code:

```
import com.liferay.portal.kernel.model.User;
import com.liferay.portal.kernel.service.UserLocalServiceUtil;
import java.util.List;

int userCount = UserLocalServiceUtil.getUsersCount();
List<User> users = UserLocalServiceUtil.getUsers(0, userCount);

for (User user:users) {
    System.out.println("User Name: " + user.getFullName());
}
```

You can even make the code somewhat Groovier:

```
import com.liferay.portal.kernel.service.UserLocalServiceUtil

userCount = UserLocalServiceUtil.getUsersCount()
users = UserLocalServiceUtil.getUsers(0, userCount)
for (user in users){
    System.out.println("User Name: " + user.getFullName())
}
```

Groovy scripts that invoke Liferay services are easy to write and execute in the script console. How to run scripts is next.

Related Topics

Running Scripts From the Script Console
Leveraging the Script Engine in Workflow
Script Examples

77.2 Running Scripts From the Script Console

The script console provides a single view for executing Groovy scripts and printing their output. It has predefined variables that facilitate printing output and working with portlets and users. Here you'll learn these things:

- How to execute a script in the script console
- The predefined variables available in the script console
- Tips for running scripts in the script console

Important: The script console is for system operations and maintenance and not for end users. Limit script console access to portal administrators.

Start with running the script console's sample script.

Running the Sample Script

Here's how to run the sample script in the script console:

1. Sign in as an administrator.
2. In the Product Menu, navigate to *Control Panel* → *Configuration* → *Server Administration*.
3. Click on *Script*. This is the script console. The default sample script prints the User count to the console output.

```
// ### Groovy Sample ###

number = com.liferay.portal.kernel.service.UserLocalServiceUtil.getUsersCount();

out.println(number);
```

4. Click *Execute* and check the script console *Output* for the User count.

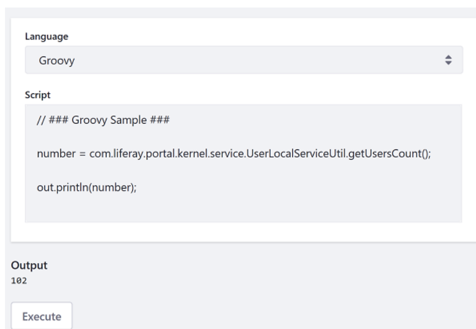

Figure 77.1: The script console's sample Groovy script prints the User count to the console's *Output* section.

The Groovy sample invokes the Liferay service utility `UserLocalServiceUtil` to get the user count. Then it uses `out` (a built-in `PrintWriter`) to write the count to the script console. Note that if you use `System.out.println` instead of `out.println`, your output is printed to Liferay's log file rather than to the script console.

Predefined Variables

Here are the predefined variables available to scripts executed in the script console:

- `out` (`java.io.PrintWriter`)
- `actionRequest` (`javax.portlet.ActionRequest`)
- `actionResponse` (`javax.portlet.ActionReponse`)
- `portletConfig` (`javax.portlet.PortletConfig`)
- `portletContext` (`javax.portlet.PortletContext`)
- `preferences` (`javax.portlet.PortletPreferences`)
- `userInfo` (`java.util.Map<String, String>`)

This script demonstrates using the `actionRequest` variable to get the portal instance's `Company`:

```
import com.liferay.portal.kernel.util.*

company = PortalUtil.getCompany(actionRequest)
out.println("Current Company:${company.getName()}\n")

out.println("User Info:")
userInfo.each {
        k,v -> out.println("${k}:${v}")
}
```

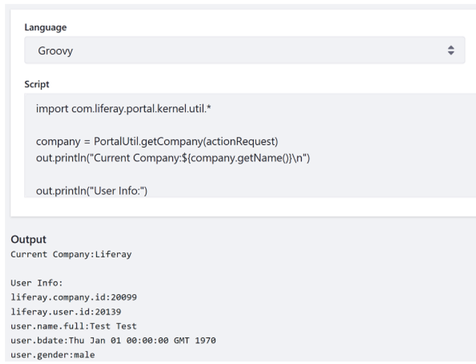

Figure 77.2: Here's an example of invoking a Groovy script that uses the predefined `out`, `actionRequest`, and `userInfo` variables to print information about the company and current user.

Tips

Keep these things in mind when using the script console:

- There is no undo.
- There is no preview.
- Permissions checking is not enforced for local services.
- Scripts are executed synchronously. Avoid executing scripts that might take a long time.

For these reasons, use the script console cautiously. Test your scripts on non-production systems before running them on production.

Of course, Liferay's script engine can be used outside of the script console. Next, you'll learn how workflows leverage Liferay's script engine.

77.3 Leveraging the Script Engine in Workflow

Liferay's Kaleo workflow engine provides a robust system for reviewing and approving content in an enterprise environment. Even if you don't leverage scripts, it's a powerful and robust workflow solution. Adding scripts takes it to the next level. These scripts aren't run from the script console, but are embedded in XML workflow definitions and run during the execution of the workflow.

The final step in a workflow runs a script that makes content available for use. The snippet below accesses the Java class associated with the workflow to set content's status to *approved*.

```
<script>
    <![CDATA[
```

```
        import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
        import com.liferay.portal.kernel.workflow.WorkflowConstants;

        WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("approved"), workflowContext);
    ]]>
</script>
<script-language>groovy</script-language>
```

At virtually any point in a workflow, you can use Liferay's script engine to access workflow APIs or other Liferay APIs. There are a lot of different ways you could use this. Here are a few practical examples:

- Getting a list of users with a specific workflow-related role
- Sending an email to the designated content approver with a list of people to contact if he is unable to review the content
- Creating an alert to be displayed in the Alerts portlet for any user assigned to approve content

Of course, before you try any of this, you need to know the appropriate syntax for inserting a script into a workflow. In an XML workflow definition, a script can be used in any XML type that can contain an `actions` tag: those types are `<state>`, `<task>`, `<fork>` and `<join>`. Inside of one of those types, format your script like this:

```
<actions>
    <action>
        <script>
            <![CDATA[*the contents of your script*]]>
        </script>
        <script-language>*your scripting language of choice*</script-language>
    </action>
    ...
</actions>
```

Here's an example of a workflow script created in Groovy. This one is used with a `Condition` statement in Kaleo. It accesses Liferay's asset framework to determine the category of an asset. The script uses the category to determine the correct approval process automatically. If the category `legal` has been applied to the asset, the asset is sent to the `Legal Review` task upon submission. Otherwise, the asset is sent to the `Default Review` task.

```
<script>
    <![CDATA[
        import com.liferay.portal.kernel.util.GetterUtil;
        import com.liferay.portal.kernel.workflow.WorkflowConstants;
        import com.liferay.portal.kernel.workflow.WorkflowHandler;
        import com.liferay.portal.kernel.workflow.WorkflowHandlerRegistryUtil;
        import com.liferay.asset.kernel.model.AssetCategory;
        import com.liferay.asset.kernel.model.AssetEntry;
        import com.liferay.asset.kernel.model.AssetRenderer;
        import com.liferay.asset.kernel.model.AssetRendererFactory;
        import com.liferay.asset.kernel.service.AssetEntryLocalServiceUtil;

        import java.util.List;

        String className = (String)workflowContext.get(
            WorkflowConstants.CONTEXT_ENTRY_CLASS_NAME);

        WorkflowHandler workflowHandler =
            WorkflowHandlerRegistryUtil.getWorkflowHandler(className);

        AssetRendererFactory assetRendererFactory =
            workflowHandler.getAssetRendererFactory();
```

```
        long classPK =
            GetterUtil.getLong((String)workflowContext.get
            (WorkflowConstants.CONTEXT_ENTRY_CLASS_PK));

        AssetRenderer assetRenderer =
            workflowHandler.getAssetRenderer(classPK);

        AssetEntry assetEntry = assetRendererFactory.getAssetEntry(
            assetRendererFactory.getClassName(), assetRenderer.getClassPK());

        List<AssetCategory> assetCategories = assetEntry.getCategories();

        returnValue = "Default Review";

        for (AssetCategory assetCategory : assetCategories) {
            String categoryName = assetCategory.getName();

            if (categoryName.equals("legal")) {
                returnValue = "Legal Review";

                return;
            }
        }
       ]]>
</script>
<script-language>groovy</script-language>
```

Within a workflow, the next task or state is chosen based on the return value. See some examples of workflow scripts by accessing the embedded workflows and inspecting the XML.

Calling OSGi Services

How do you call OSGi services from a workflow script, accounting for the dynamic environment of the OSGi runtime, where services your script depends on can disappear without notice? Use a service tracker. That way you can make sure your code has access to the service it needs, and if not, do something appropriate in response. Here's a little example code to show you how this might look in Groovy:

```
import com.liferay.journal.model.JournalArticle;
import com.liferay.journal.service.JournalArticleLocalService;
import com.liferay.portal.scripting.groovy.internal.GroovyExecutor;

import org.osgi.framework.Bundle;
import org.osgi.framework.FrameworkUtil;
import org.osgi.util.tracker.ServiceTracker;

ServiceTracker<JournalArticleLocalService, JournalArticleLocalService> st;

try {
    Bundle bundle = FrameworkUtil.getBundle(GroovyExecutor.class);

    st = new ServiceTracker(bundle.getBundleContext(), JournalArticleLocalService.class, null);
    st.open();

    JournalArticleLocalService jaService = st.waitForService(500);

    if (jaService == null) {
        _log.warn("The required service 'JournalArticleLocalService' is not available.");
    }
    else {
        java.util.List<JournalArticle>articles = jaService.getArticles();
        if (articles != null) {
            _log.info("Article count: " + articles.size());
```

```
        } else {
            _log.info("no articles");
        }
    }
}
catch(Exception e) {
    //Handle error appropriately
}
finally {
    if (st != null) {
        st.close();
    }
}
```

If you read the article on service trackers, the only odd looking piece of the above code is the `getBundle` call: why is `GroovyExecutor.class` passed as a parameter? The parameter passed to the `FrameworkUtil.getBundle` call must be a class from the OSGi bundle executing the workflow script. This is different from the context of a plugin project, where you'd want to get the bundle hosting the class where you're making the call (using `this.getClass()`, for example). Note that for another scripting engine, you must pass in a concrete class from the particular bundle executing your script.

The combination of Liferay's script and workflow engines is incredibly powerful. Since, however, it enables users to execute code, it can be dangerous. When configuring your permissions, be aware of the potential consequences of poorly or maliciously written scripts inside a workflow definition. For more information on creating workflow definitions with Kaleo workflow, see the workflow documentation.

Related Topics

Running Scripts From the Script Console

Script Examples

77.4 Script Examples

Here are some examples to help you use Liferay's script console. Note: Most of these originated from a Liferay blog post.

The following scripts are Groovy scripts but they can be adapted to other languages.

- Example 1: Presenting New Terms of Use to Users
- Example 2: Embedding HTML Markup in Script Outputs
- Example 3: Show Exceptions in the Script Console
- Example 4: Implement a Preview Mode
- Example 5: Plan a File Output for Long-Running Scripts

Example 1: Presenting New Terms of Use to Users

This example retrieves user information from the database, makes changes, and then saves the changes in the database. Suppose that your company has updated the terms of use and wants present users with the updated terms of use whenever they sign in next. When they agree to the terms of use, a boolean attribute called `agreedToTermsOfUse` is set in their user records. As long as the value of this variable is `true`, users aren't presented with the terms of use when they sign in. If you set this flag to `false` for each user, each user must agree to the terms of use again before they can sign in.

1. Enter and execute the following code in the script console:

    ```
    import com.liferay.portal.kernel.service.UserLocalServiceUtil

    userCount = UserLocalServiceUtil.getUsersCount()
    users = UserLocalServiceUtil.getUsers(0, userCount)

    long currentUserId = Long.parseLong(userInfo.get("liferay.user.id"))

    for (user in users) { println("User Name: " + user.getFullName() + " -- " +
    user.getAgreedToTermsOfUse()) }
    ```

    This code prints each user's `agreedToTermsOfUse` attribute value.

2. Replace that with this script:

    ```
    import com.liferay.portal.kernel.service.UserLocalServiceUtil

    userCount = UserLocalServiceUtil.getUsersCount()
    users = UserLocalServiceUtil.getUsers(0, userCount)

    for (user in users){

        if(!user.isDefaultUser() && (user.getUserId() != currentUserId)) {

                user.setAgreedToTermsOfUse(false)
                UserLocalServiceUtil.updateUser(user)

        }

    }
    ```

    This sets each user's `agreedToTermsOfUse` attribute to `false`. It skips the default user as well as the default admin user that's currently signed in and running the script.

3. Click *Execute*.

4. Verify the script updated the records by running the first script again.

    All users (except the default user and your user) have been updated.

You've enabled the new terms of use agreement for all users to accept.

Example 2: Embedding HTML Markup in Script Outputs

The output of the script console is rendered as HTML content. Thus, you can embed HTML markup in your output to change its look and feel. Here's an example:

```
import com.liferay.portal.kernel.service.*

number = com.liferay.portal.kernel.service.UserLocalServiceUtil.getUsersCount();
out.println(
        """
                <div style="background-color:black; text-align: center">
                        <h1 style="color: #37A9CC; font-size:xx-large">${number}</h1>
                </div>
        """);
```

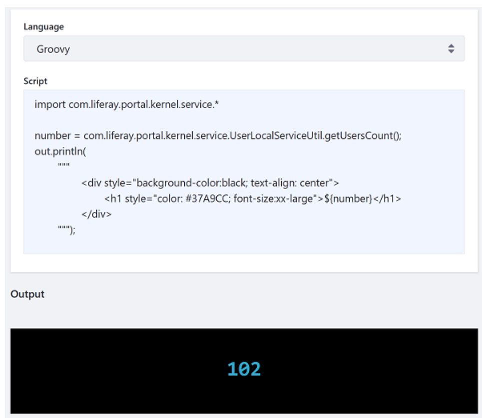

Figure 77.3: Here's an example of invoking a Groovy script that embeds HTML markup in the output of the script.

Example 3: Show Exceptions in the Script Console

When any exception occurs during script execution, the error message is always the same:

`Your request failed to complete.`

This message gives no detail about the error. To find information about the error and what caused it, you must usually examine the server logs.

You can, however, use the following technique to make exception details appear in the script console. Wrap your code with a try / catch block and print the stacktrace to the console output from the catch clause. Note that even this technique does not catch script syntax errors. Here's an example:

```
try {
        nullVar = null
        out.println(nullVar.length())
} catch(e) {
        out.println("""<div class="portlet-msg-error">${e}</div>""")
        e.printStackTrace(out)
}
```

Figure 77.4: Here's an example of a Groovy script that catches exceptions and prints exception information to the script console.

Example 4: Implement a Preview Mode

Since Liferay's script console does not provide an undo feature, it can be convenient to set up a kind of preview mode. The purpose of a preview mode is to determine any permanent effects of a script before any information is actually saved to the Liferay database. The preview mode consists in using a `previewMode` flag which determines whether the operations with permanent effects should be executed or not. If `previewMode` is `true`, all the data that would be permanently affected by the script is printed instead. Then you can see an outline of the data impacted by the script. If everything is okay, switch the flag so the script can make permanent updates to the database.

Here's an example Groovy script that sets users to inactive. Clearly, you'd want to test this with preview mode before running it:

```
import java.util.Calendar
import com.liferay.portal.kernel.service.*
import com.liferay.portal.kernel.model.*
import com.liferay.portal.kernel.dao.orm.*
import static com.liferay.portal.kernel.workflow.WorkflowConstants.*

//
// Deactivate users never logged and created since more than 2 years
//

previewMode = true // Update this flag to false to really make changes

Calendar twoYearsAgo = Calendar.getInstance()
twoYearsAgo.setTime(new Date())
twoYearsAgo.add(Calendar.YEAR, -2)

DynamicQuery query = DynamicQueryFactoryUtil.forClass(User.class)
        .add(PropertyFactoryUtil.forName("lastLoginDate").isNull())
        .add(PropertyFactoryUtil.forName("createDate").lt(twoYearsAgo.getTime()))

users = UserLocalServiceUtil.dynamicQuery(query)

users.each { u ->
         if(!u.isDefaultUser() && u.getStatus() != STATUS_INACTIVE) {
                out.println(u.getEmailAddress())
                if(!previewMode) {
                        UserLocalServiceUtil.updateStatus(u.getUserId(), STATUS_INACTIVE)
                }
         }
}
```

```
if(previewMode) {
        out.println('Preview mode is on: switch off the flag and execute '
                + 'again this script to make changes to the database')
}
```

Example 5: Plan a File Output for Long-Running Scripts

Important: The script console is for system operations and maintenance and not for end users. Limit script console access to portal administrators.

When a script has been running for a long time, the console could return an error even though the script can continue running and potentially conclude successfully. But it's impossible to know the outcome without the corresponding output!

To bypass this limitation, you can send the output of the script console to a file instead of to the console itself or to the Liferay log. For example, consider this script:

```
import com.liferay.portal.kernel.service.*
import com.liferay.portal.kernel.dao.orm.*

// Output management

final def SCRIPT_ID = "MYSCRIPT"
outputFile = new File("""${System.getProperty("liferay.home")}/scripting/out-${SCRIPT_ID}.txt""")
outputFile.getParentFile().mkdirs()

def trace(message) {
        out.println(message)
        outputFile << "${message}\n"
}

// Main code

users = UserLocalServiceUtil.getUsers(QueryUtil.ALL_POS, QueryUtil.ALL_POS)
users.each { u ->
        trace(u.getFullName())
}
```

The script above creates a subfolder of Liferay Home called `scripting` and saves the script output to a file in this folder. After running the script above, you can read the generated file without direct access to the file system. Here's a second script that demonstrates this:

```
final def SCRIPT_ID = "MYSCRIPT"
outputFile = new File("""${System.getProperty("liferay.home")}/scripting/out-${SCRIPT_ID}.txt""")
out.println(outputFile.text)
```

One advantage of using a dedicated output file instead of using a classic logger is that it's easier to get the script output data back. Getting the script output data would be more difficult to obtain from the portal log, for example, because of all the other information there.

Related Topics

Running Scripts From the Script Console

Leveraging the Script Engine in Workflow

Using Liferay's Script Engine

77.5 Custom Fields

Have you ever wondered why there's no *Head Circumference* field in the form for adding users to Liferay DXP? Probably because most sites based on Liferay DXP don't need it. If you're an administrator at the Lunar Resort, however, you certainly need to know your guests' head circumference so you can provide them with a properly fitting helmet.

Many of Liferay DXP's assets and resources let you add new fields to their edit forms. Here's the complete list:

- Blogs Entry
- Bookmarks Entry
- Bookmarks Folder
- Calendar Booking
- Document
- Documents Folder
- Message Boards Category
- Message Boards Message
- Organization
- Page
- Role
- Site
- User
- User Group
- Web Content Article
- Web Content Folder
- Wiki Page

Developer Use Case: Adding custom fields to Liferay DXP resources affords flexibility to developers. Suppose you must limit the number of users that can be assigned to a particular Role. First an administrator creates a custom field called *max-users* for the Role. A developer then creates a module that inserts logic before a user is added to that Role. If the logic detects that the maximum number of Role users would be exceeded by completing the action, an exception is thrown and the action does not complete.

Adding Custom Fields

To add custom fields, find the Custom Fields entry beneath the Control Panel's Configuration heading.

To add a custom field to one of the listed entities,

1. Choose a resource by clicking on it.

2. Click the add (+) button.

3. Add a Key.

```
**Note:** The Key you enter here is the name of the new field. It's stored
in the database and used by developers to access the custom field with the
`<liferay-ui:custom-attribute />` tag. It is also used to label the
field in the UI.
```

---

4. Choose the type of field you want to add.

```
**Available Field Types:** There are a number of field types to choose from:

Presets:

- Selection of Integer Values
- Selection of Decimal Values
- Selection of Text Values
- Text Box
- Text Box--Indexed
- Text Field--Secret
- Text Field--Indexed

Primitives

- True/False
- Date
- Decimal number (64-bit)
- Group of Decimal numbers (64-bit)
- Decimal number (32-bit)
- Group of Decimal numbers (32-bit)
- Integer (32-bit)
- Group of Integers (32-bit)
- Integer (64-bit)
- Group of Integers (64-bit)
- Decimal Number or Integer (64-bit)
- Group of Decimal numbers or Integer (64-bit)
- Integer (16-bit)
- Group of Integers (16-bit)
- Text
- Group of Text Values
- Localized Text
```

```
![ At The Lunar Resort, a Head Circumference field is necessary for all users.](./images/custom-fields-user-head-circumference.png)
```

5. Click Save.

That's it.

Once you have a custom field for a resource, go find it in the UI of the resource. First find the UI location for the resource, and all its custom fields are displayed in a Custom Fields panel. For example, consider the Users UI:

Navigate to *Control Panel* → *Users and Organizations*. Click on a User to open the Edit User form and scroll down to find your custom field.

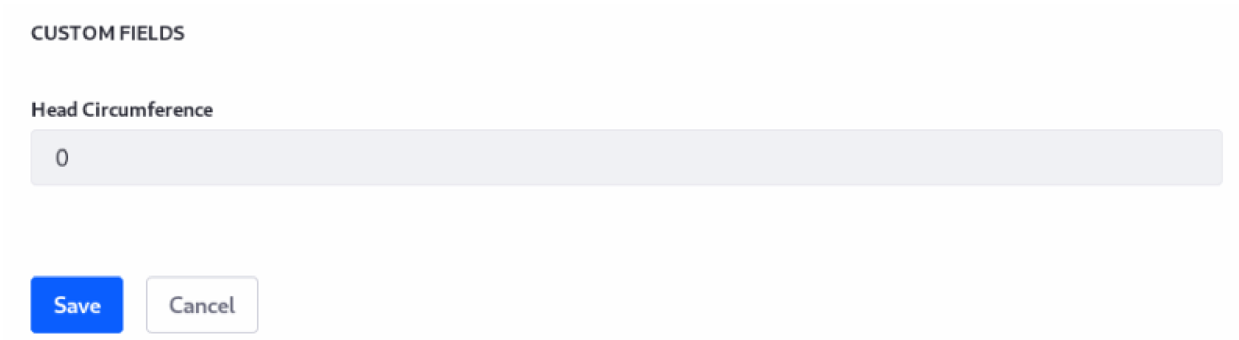

Figure 77.5: The Custom Fields panel is found at the bottom of the Edit User form.

Figure 77.6: Custom Fields can be configured after you create them. The exact options depend on the field type you choose.

Editing a Custom Field

You can't change the key or field type of a custom field, but you can delete it and create a new one if necessary. Additionally, once you create the custom field, there are properties you can configure. For example, here's the Head Circumference field's edit page:

Edit an individual custom field's permissions by clicking the field's kebab menu (⋮), then *Permissions*. Permission can be granted or removed for these actions:

- Delete
- Permissions
- Update
- View

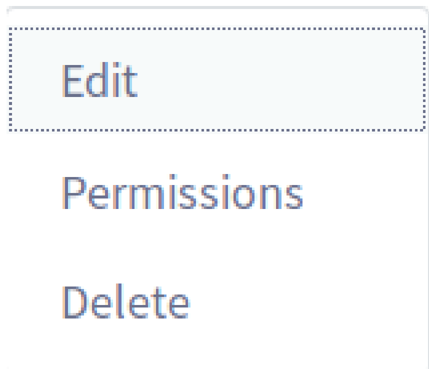

Figure 77.7: You can delete a custom field, edit it, or configure its permissions.

Custom fields make many of Liferay DXP's entities extensible directly from the administrative user interface. Use them as is or combine them with some back-end code, and you have yet another powerful, flexible feature at your disposal. As they're fond of saying at The Lunar Resort, "The sky is certainly not the limit."

77.6 Polls

How can The Lunar Resort stay connected with its earthbound clientèle from 239,000 miles away? Make them feel really involved and enthusiastic about the resort by asking them for feedback. You're not just creating a poll, you're making connections.

Use Polls to find out what your site visitors are thinking and keep them engaged with your site's content.

Two applications make and display a poll: the *Polls* application in the Site Menu and the *Polls Display* widget you add to a page.

Creating a Poll

From the Site Menu, go to *Content* → *Polls*.

1. Click the + button and fill out the form.

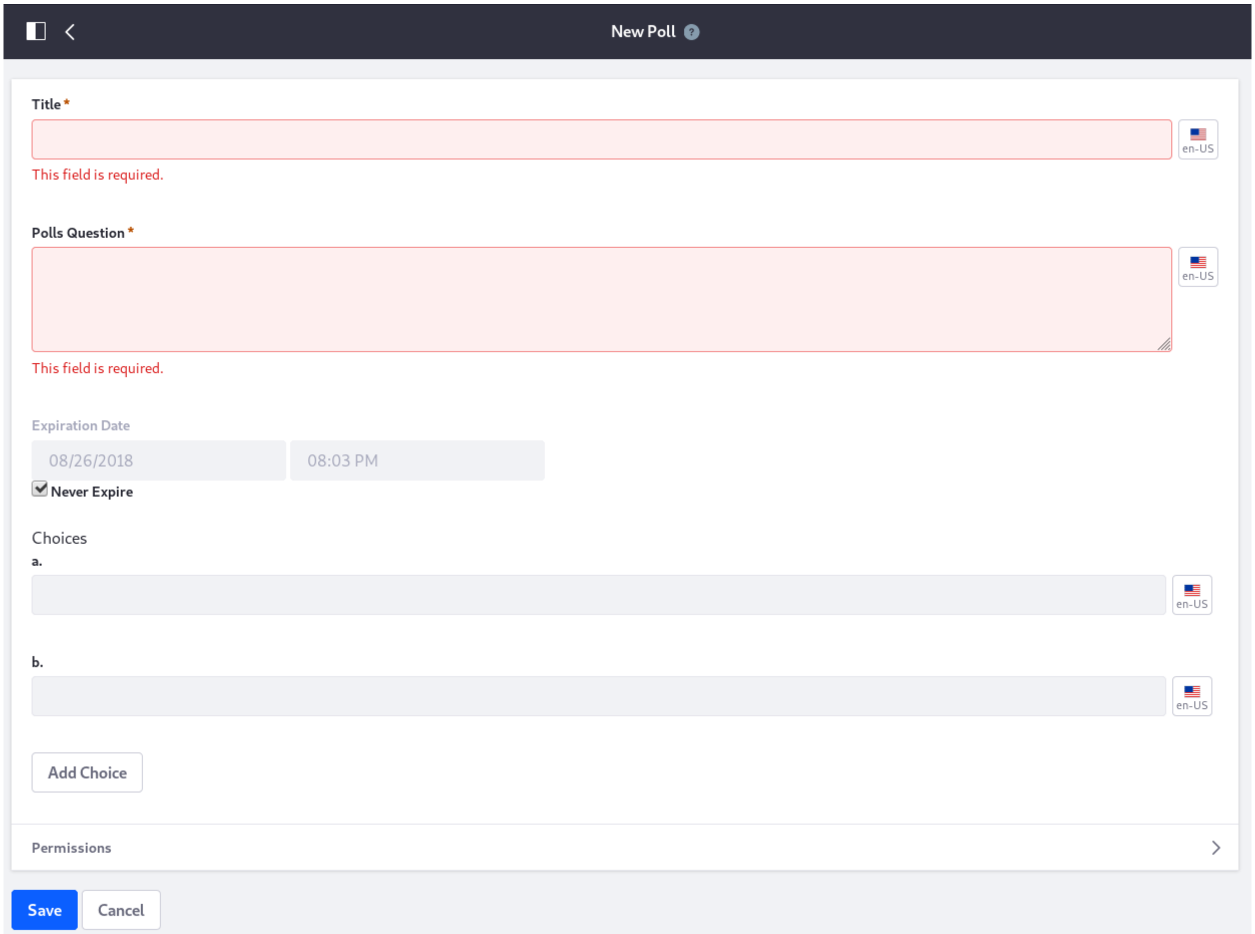

Figure 77.8: Besides the Title and the Polls Question, you must enter data for each of the Choices fields when creating a new poll.

Title: (Required) Enter the name of the poll question.

Polls Question: (Required) Enter the text of the poll question.

Expiration Date: Enter the date and time you want the poll to expire.

Choices: Enter at least two options for the poll question.

Add Choice: Enter additional answer options for the poll question.

Permissions: Manage who can view and edit the poll.

2. Click *Save* to add the poll to the Polls application.

Once a poll is created, the Polls Display portlet publishes it until it expires or is deleted. Set an expiration date for a poll by selecting the day and time in the Add Poll form. The default is set to *Never Expire*.

When a published poll expires, the poll results are displayed, but users can't add new entries. To remove an expired poll from a page, remove the Poll Display portlet or configure it to show another poll question. See the section below for more details about the Polls Display portlet.

Permissions are set on individual polls. Use permissions, for example, to allow some privileged users to vote on a certain poll question, while others can only view it.

Creating a poll is fairly straightforward. Next, complete the two-step process and put your poll on a page.

Adding a Poll to a Page

Now that you have created your poll question, you can present it to your users:

1. Go to a page and add a Polls Display widget from *Add* → *Widgets* → *Content Management*.
2. Click *Please configure this portlet to make it visible to all users.*
3. In the dialog box that appears, select the poll to display.
4. Click *Save*.

Once the poll question has been placed on the page, perform other tasks using the icons at the bottom of the portlet.

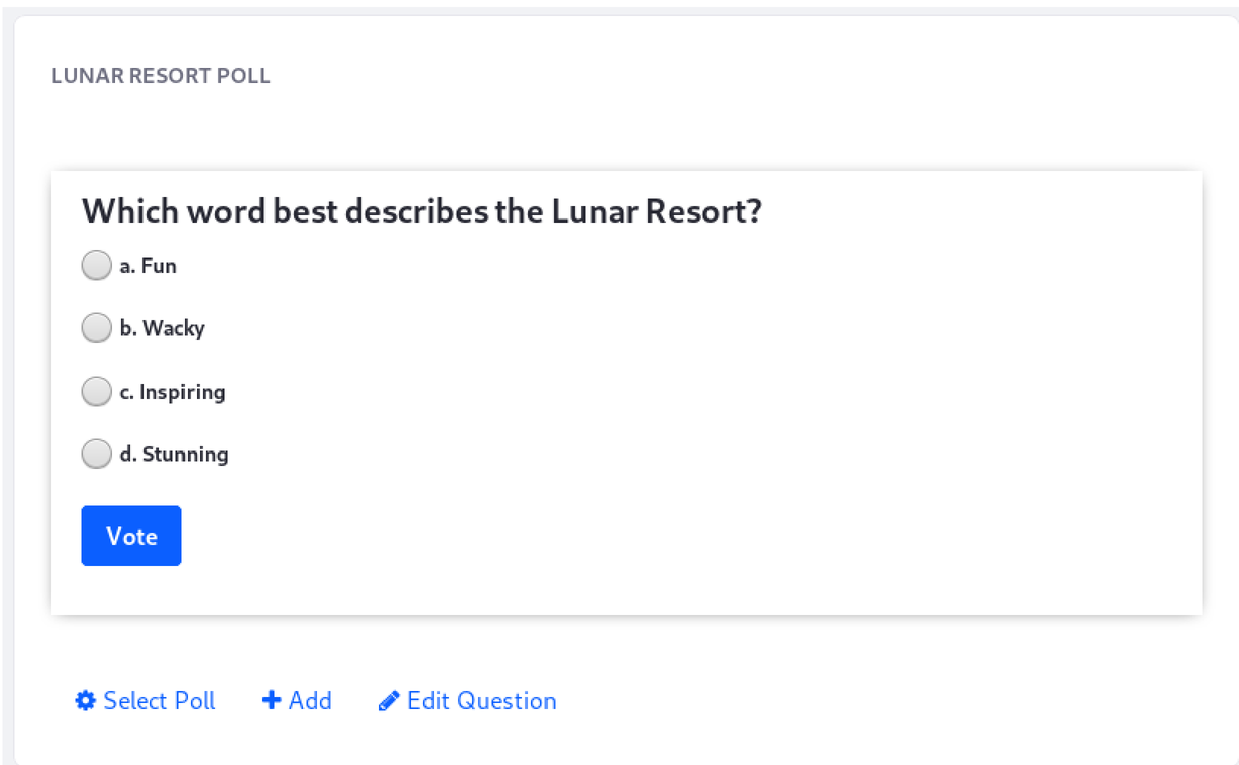

Figure 77.9: These buttons provide shortcuts to the widget's configuration, as well as to some of the Polls Application's functionality.

Edit Question: Displays a similar dialog box to the one used to create the poll.

Select Poll: Displays the same dialog box as Configuration, allowing you to choose different polls from the drop-down menu.

Add: Allows you to create a new poll.

Viewing Poll Results

All the polls you create appear in the Polls portlet in the Site Menu → *Content* → *Polls*. When users vote in the poll, the data is collected here. Click on a poll to see a breakdown of the results.

Figure 77.10: Selecting a poll in the Polls portlet puts the data at your fingertips.

If you click on one of the listed *Charts*, the portlet generates an appropriate visualization of the data.

Below this is an item called *Charts*. This option shows the poll results represented in various graphs. The graphs are *Area*, *Horizontal Bar*, *Line*, *Pie*, and *Vertical Bar*.

There is also a listing of the users who voted in the poll, how they voted, and a time/date stamp of when their votes were cast. Registered users are represented by name. Guest users have a blank *User* field.

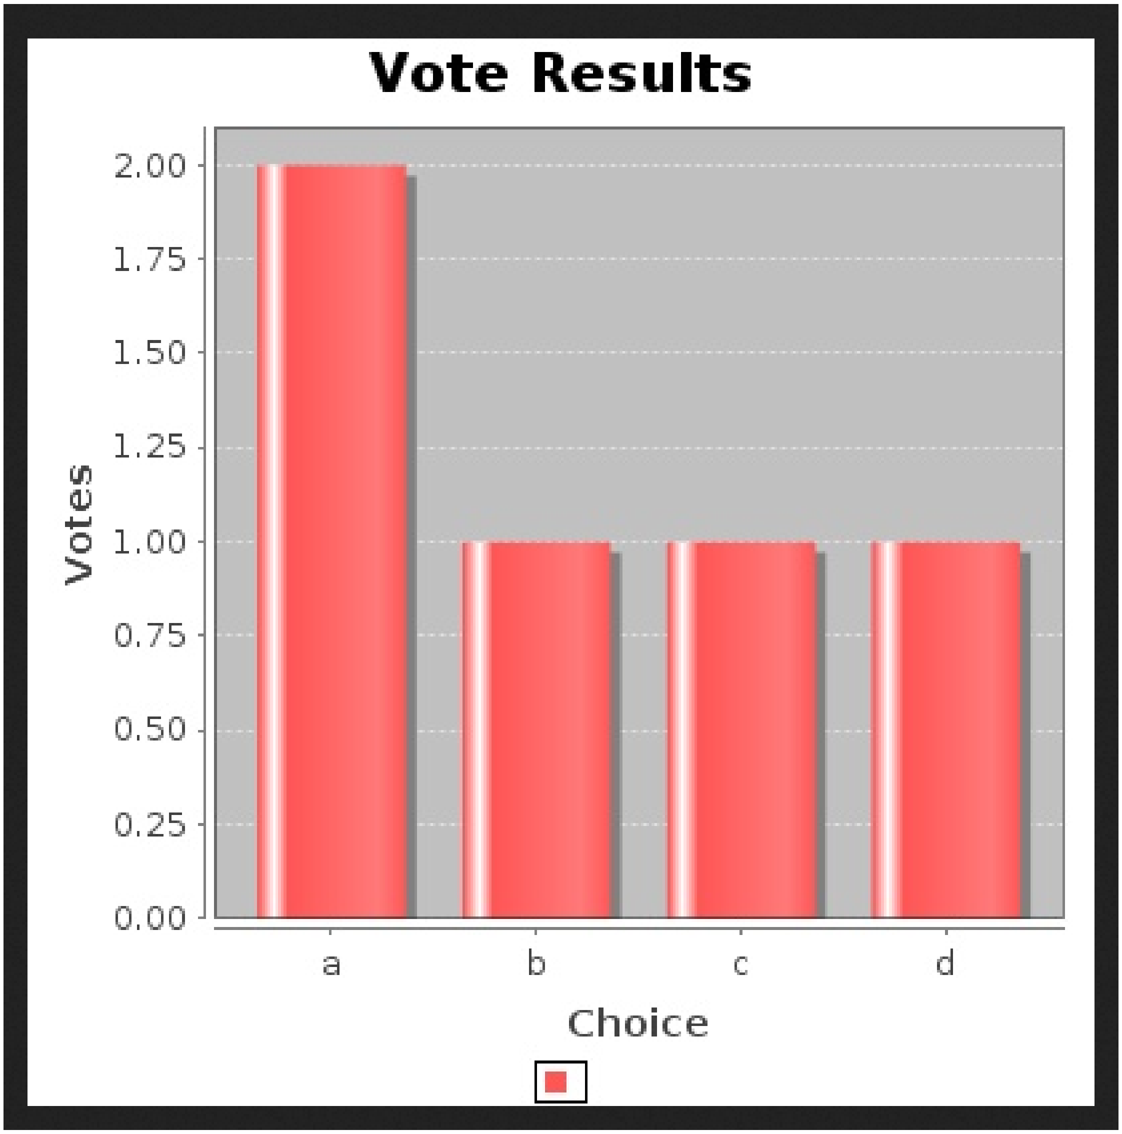

Figure 77.11: This is what the vertical bar graph for the Lunar Resort poll results looks like.

CHAPTER 78

USING THE CALENDAR

The Calendar widget is an updated, digitized, 3D-printed sundial. Okay, it's really a tool for storing and sharing scheduled events. It's a personal planner for individual users, a shared calendar for an entire site, or both at the same time. It can be used to create multiple calendars for a single Site or User, to overlay the events stored in multiple calendars for simultaneous view, to send email reminders to users, and more.

Note: The calendar supports social activities. Whenever a calendar event is added or updated, a corresponding social activity notification is created. If the event was added or updated in a calendar that the current user has permission to view, the social activity is viewable in the Activities widget.

78.1 Configuring the Calendar

Once the Calendar widget is on a page, open the  menu in the widget's header and click *Configuration*.

From the *User Settings* tab, customize the calendar's default view and settings.

Time Format: Choose *Locale*, *AM/PM*, or *24 Hour*. *Locale* is a dynamic setting that chooses whether to display the time in *AM/PM* or *24 Hour* format, based on the preferences set by the User's locale. *AM/PM* displays times such as 8AM or 11PM. The *24 Hour* time format displays times such as 08:00 and 23:00.

Default Duration: Choose an event duration. When you add a new event to the calendar, the time you set here specifies how long events last by default.

Default View: Choose *Day*, *Week*, *Month* or *Agenda*. This sets the default for when the calendar is first displayed, but the view can be changed by clicking the appropriate button at the top-right of the widget.

Week Starts On: Choose *Sunday*, *Monday*, or *Saturday*.

Time Zone: Choose a time zone or check the *Use Global Time Zone* box.

If you check *Use Global Time Zone*, the time displayed depends on whether it's being viewed by a logged-in user or a guest. If a user is logged in, the Calendar displays events using the time zone set for the user in *User Menu* → *My Account* → *Account Settings* → *Preferences* → *Display Settings* →

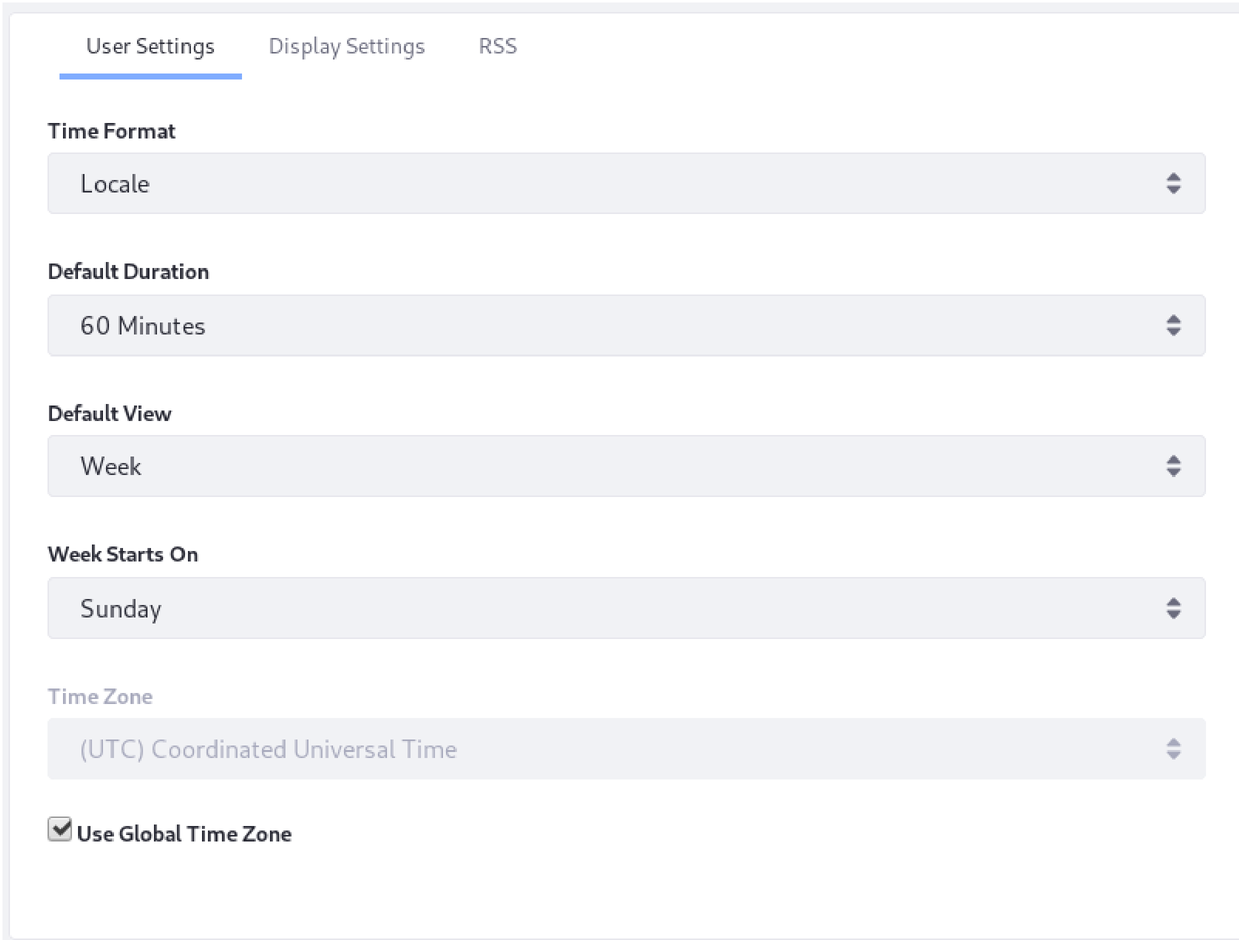

Figure 78.1: The Setup → User Settings tab provides the options you need to get started quickly.

Time Zone. If the Calendar is viewed by a guest or a user who is not logged in, the Calendar displays events using the time zone set by the portal administrator in *Control Panel* → *Portal Settings* → *Display Settings* → *Time Zone*.

From the *Display Settings* tab, set the display behavior for the calendar.

Display Scheduler Only: By default, the list of calendars and a mini-calendar view (used for quickly navigating to a particular date) are displayed. Check this to display only the scheduler (the large calendar view showing the calendar and scheduled events).

Display User Events: Turns off the display of the current, logged in User's personal calendar and events.

Display Scheduler's Header: If disabled, removes the ability to toggle through the calendar views (for example, Day/Week/Month/Agenda) and access to the Add Event button.

Enabled Views: If one of the available views is disabled (Day, Week, Month, Agenda), it disappears from the scheduler's header.

Maximum Days to Display: Set the maximum number of days to display in the Agenda view.

Maximum Events to Display: Set the maximum number of events to display in the Agenda view.

Use the *RSS* tab to disable RSS subscription or configure the RSS behavior.

Enough with configuration. Next you'll learn how to use it.

78.2 Using the Calendar Widget

The calendar widget displays a small monthly calendar showing an overview of upcoming events. A larger area shows the Scheduler, a more detailed calendar with a number of options: you can set it to to display a *Day*, *Week*, or *Month*, or choose a more event-oriented *Agenda* setting.

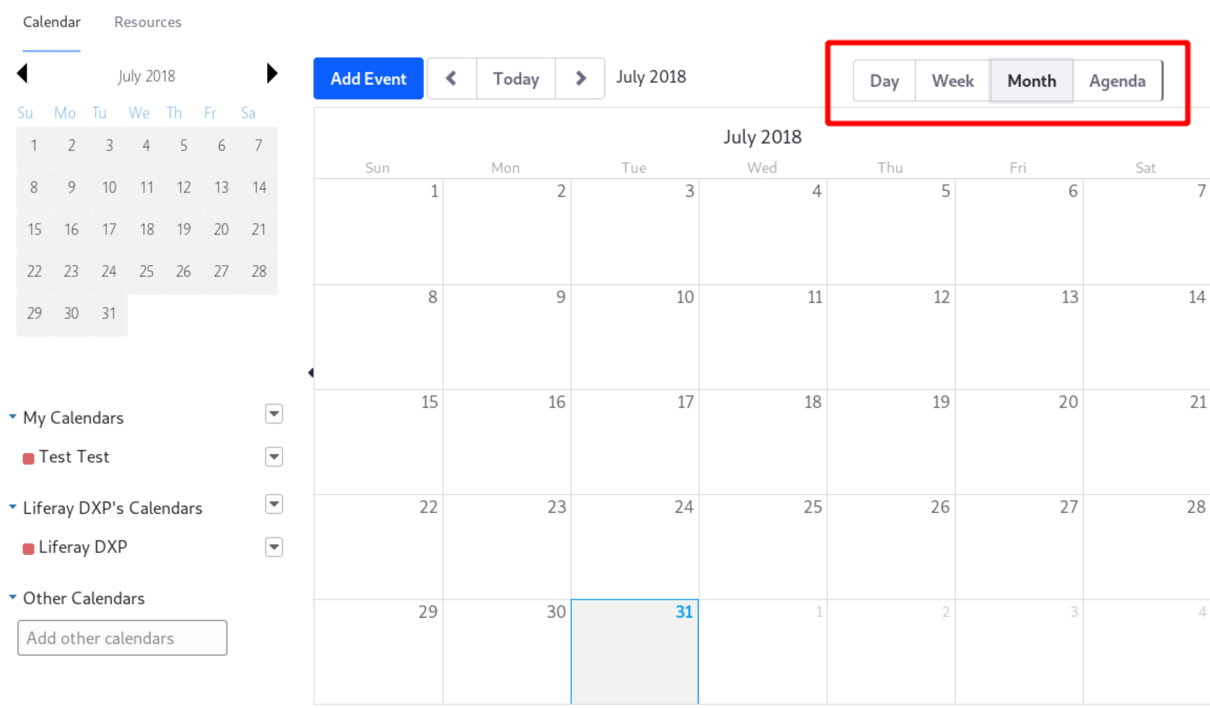

Figure 78.2: The default view is set in configuration, but a user can change it at any time.

Two calendars are included by default when the widget is first added to a page: a personal calendar for the current user and a Site calendar for the current Site. These are displayed in the widget's lower left. Next to each calendar is a colored box: click it to show/hide that calendar's events in the main viewing area.

Adding New Calendars

To create a new personal calendar,

1. Click on the arrow to the right of the *My Calendars* header and select *Add Calendar* from the menu.

2. Fill in the *Add Calendar* form. Give the calendar a name and description, set a time zone, and decide if it's your user's *default calendar*—the one that is shown automatically whenever the widget is displayed. You can also pick a color, which color codes events whenever multiple

calendar's events are displayed at once. You can also decide to enable ratings or comments on the calendar's events, and configure permissions.

To edit an existing calendar instead of adding a new one, select *Manage Calendars* from the menu.

To add or edit a Site calendar, open the menu next to the header with the Site's name.

Adding Events to a Calendar

To add events to a calendar,

1. Click on any day in the main viewing area to open an event creation pop up. If you've selected the *Day* or *Week* view, you can click on the specific time when your event begins.

2. Name your event and assign it to a calendar. Click *Save* to create the event immediately or *Edit* to enter additional information.

3. If you clicked *Edit*, complete the *Edit Event* form. Enter start and end times and enter a description. To schedule an event that reoccurs, check the *Repeat* box and fill in the *Repeat* pop up.

Additional Event Functions

At the bottom of the *Edit Event* form, there are several collapsed sections: *Details*, *Invitations*, *Reminders*, *Categorization*, and *Related Assets*.

Details In the Details section, you can move the event to another calendar, enter a location, and configure the event's permissions.

Giving a user permission to add, delete, or update discussion allows them to make, edit and remove comments on the event. The *Permissions* permission allows a Role to update permissions for the event.

Invitations In the invitations section, invite Users, Sites, or Calendar Resources (see the next tutorial for more on resources: in brief, a resource is anything you might need for an event—a conference room, a vehicle, etc.). Follow these steps:

1. Enter the name of an invitee (User, Site, or Resource) in the *Invitations* field. Hit *Enter* to add them to the *Pending* column.

2. Check the availability of invitees by clicking the arrow next to the their names and selecting *Check Availability*. This displays their calendars (assuming you have permission to view them).

An automated email is sent to invitees who must navigate to the calendar widget to respond. See below to customize the content of the invitation.

When invitees respond to the invitation, their names move to the *Accepted*, *Declined*, or *Maybe* columns.

Reminders Schedule up to two email reminders to send to attendees. Reminders translate the time of the event into the recipients own time zone.

See below to customize the content of the reminder email.

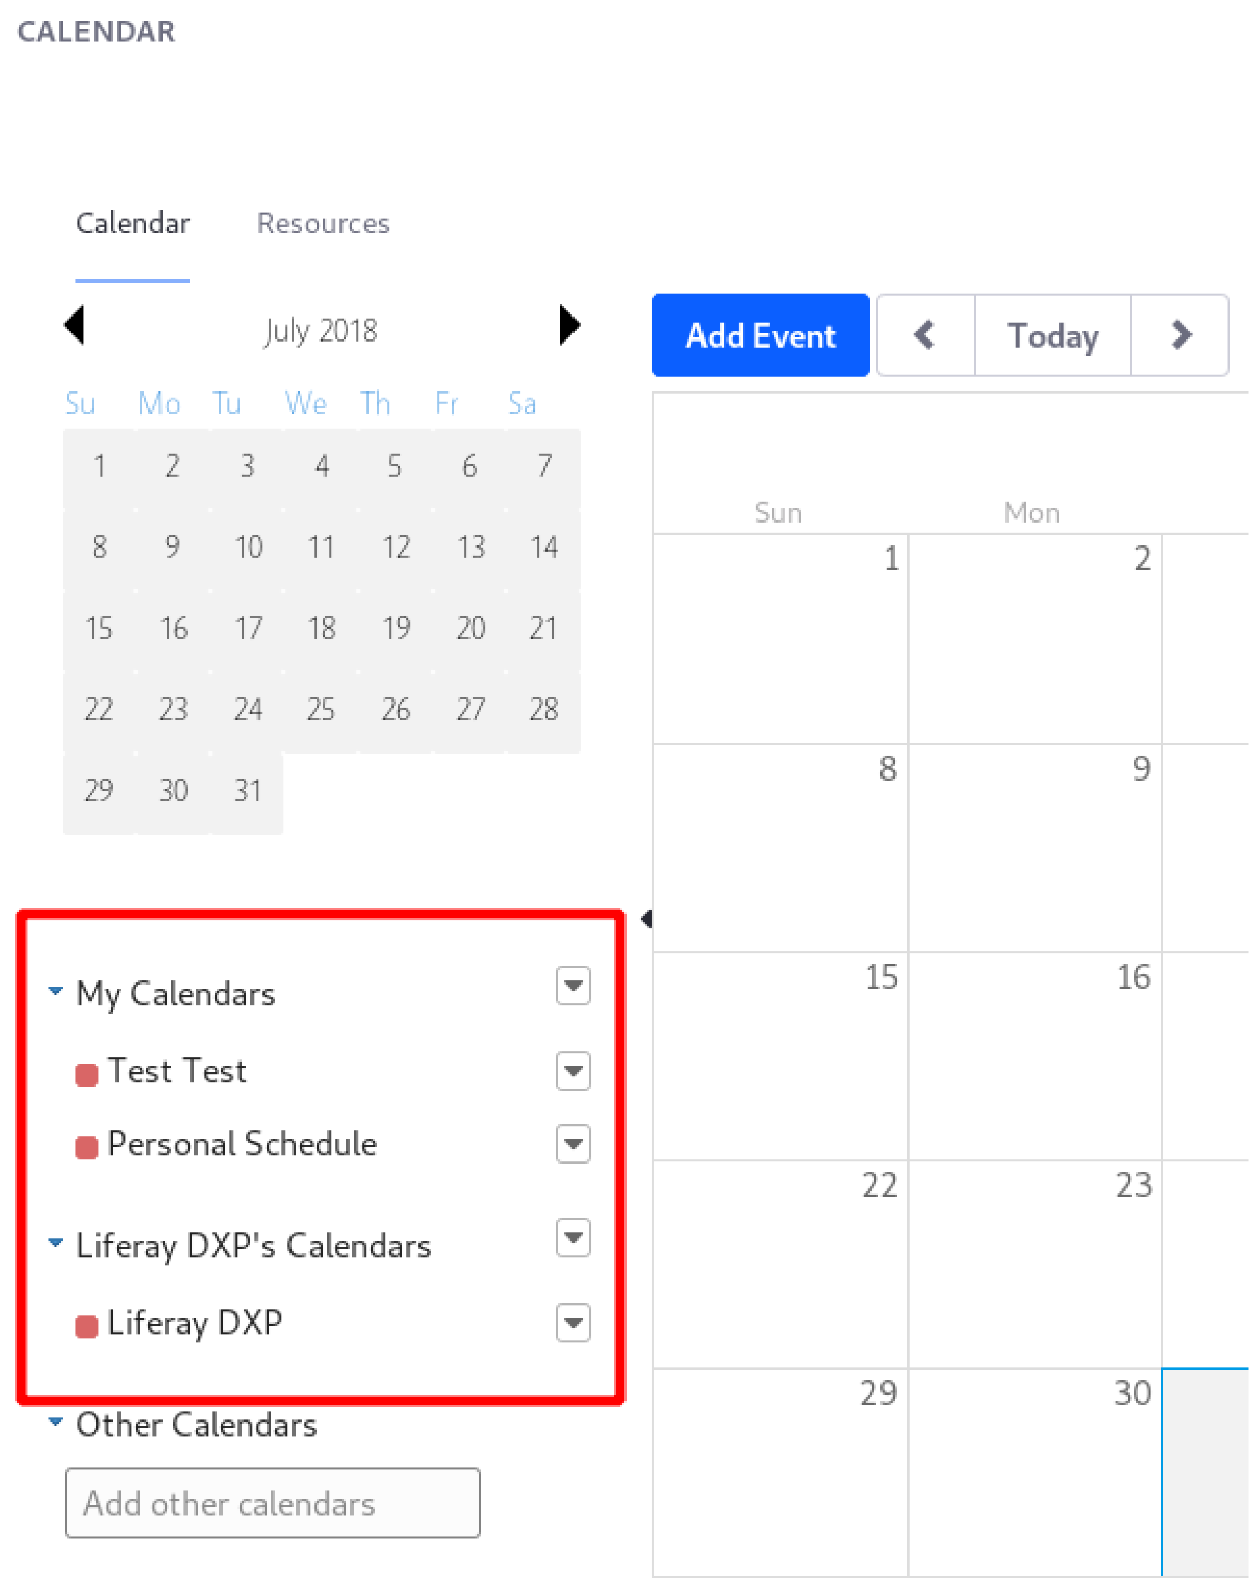

Figure 78.3: Personal and Site calendars are shown in the lower left. This image shows calendars belonging to User *Test Test* and Site *Liferay DXP*.

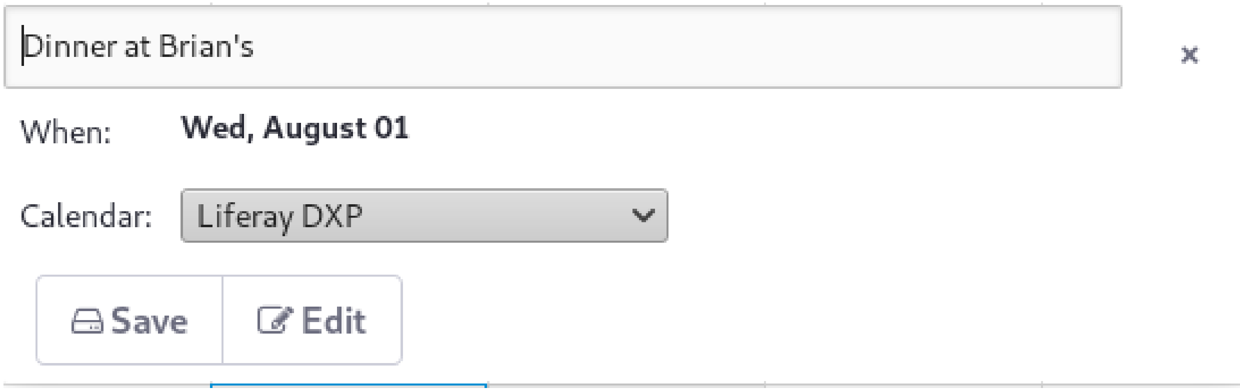

Figure 78.4: When you click anywhere on the calendar, you'll see the event creation pop up appear. Click *Edit* to specify details for your event.

Categorization Tag your event or assign it to a category so it appears in appropriate search results and is published by any asset publisher set to publish content assigned to the same category.

Related Assets List an asset—such as an agenda or supplementary material for a meeting—as related to your event. Links to related assets are displayed in the *Event Details* window.

Customizing Email Notifications

To customize email notifications for event invitations and reminders,

1. From the Calendar widget, click on the arrow next to a calendar and select *Calendar Settings*.
2. Click on the *Notification Templates* tab. Then select either the *Invite Email* or the *Reminder Email* sub-tab.
3. Edit the email as desired. At the bottom of the screen is a glossary that specifies variables for terms that were set when you created the event. Use these variables to refer to event-specific information, such as the event's name, date or location. It's a good idea to include a link to the event (use the variable `[$EVENT_URL$]`) as users must navigate to the calendar widget in order to respond.

Click *Save*. Now your notifications contain the proper text.

The next article covers setting up calendar resources and porting data from one installation to another.

78.3 Calendar Resources and Porting

With calendar resources, you can invite entities other than people to your events. This is beneficial for finding the availability of important resources your event requires, like a conference room, laptop, or, at The Lunar Resort, the Sasquatch Space Suit used to scare guests out on Lunar hikes.

Another important topic is porting your calendar's data from one installation of Liferay DXP to another.

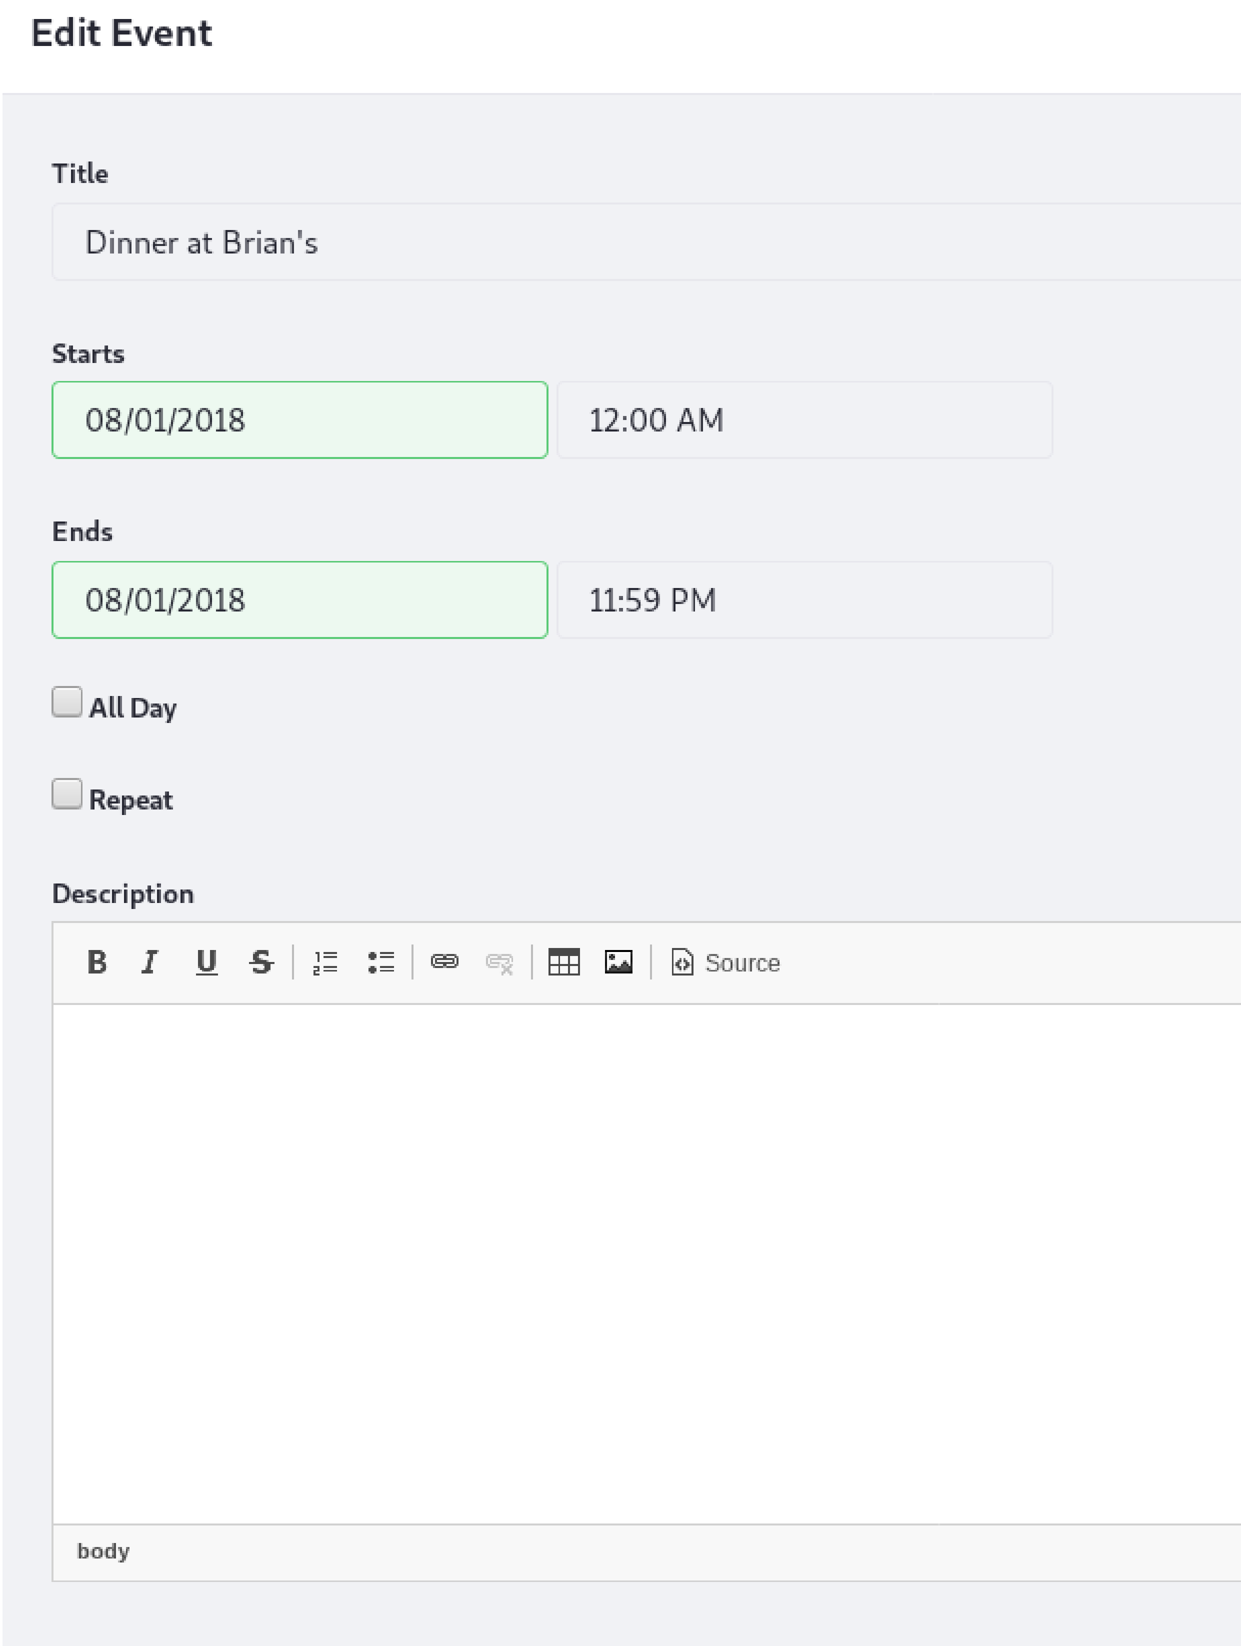

Figure 78.5: You can specify event details such as the event title, start date, end date, description, location, and more.

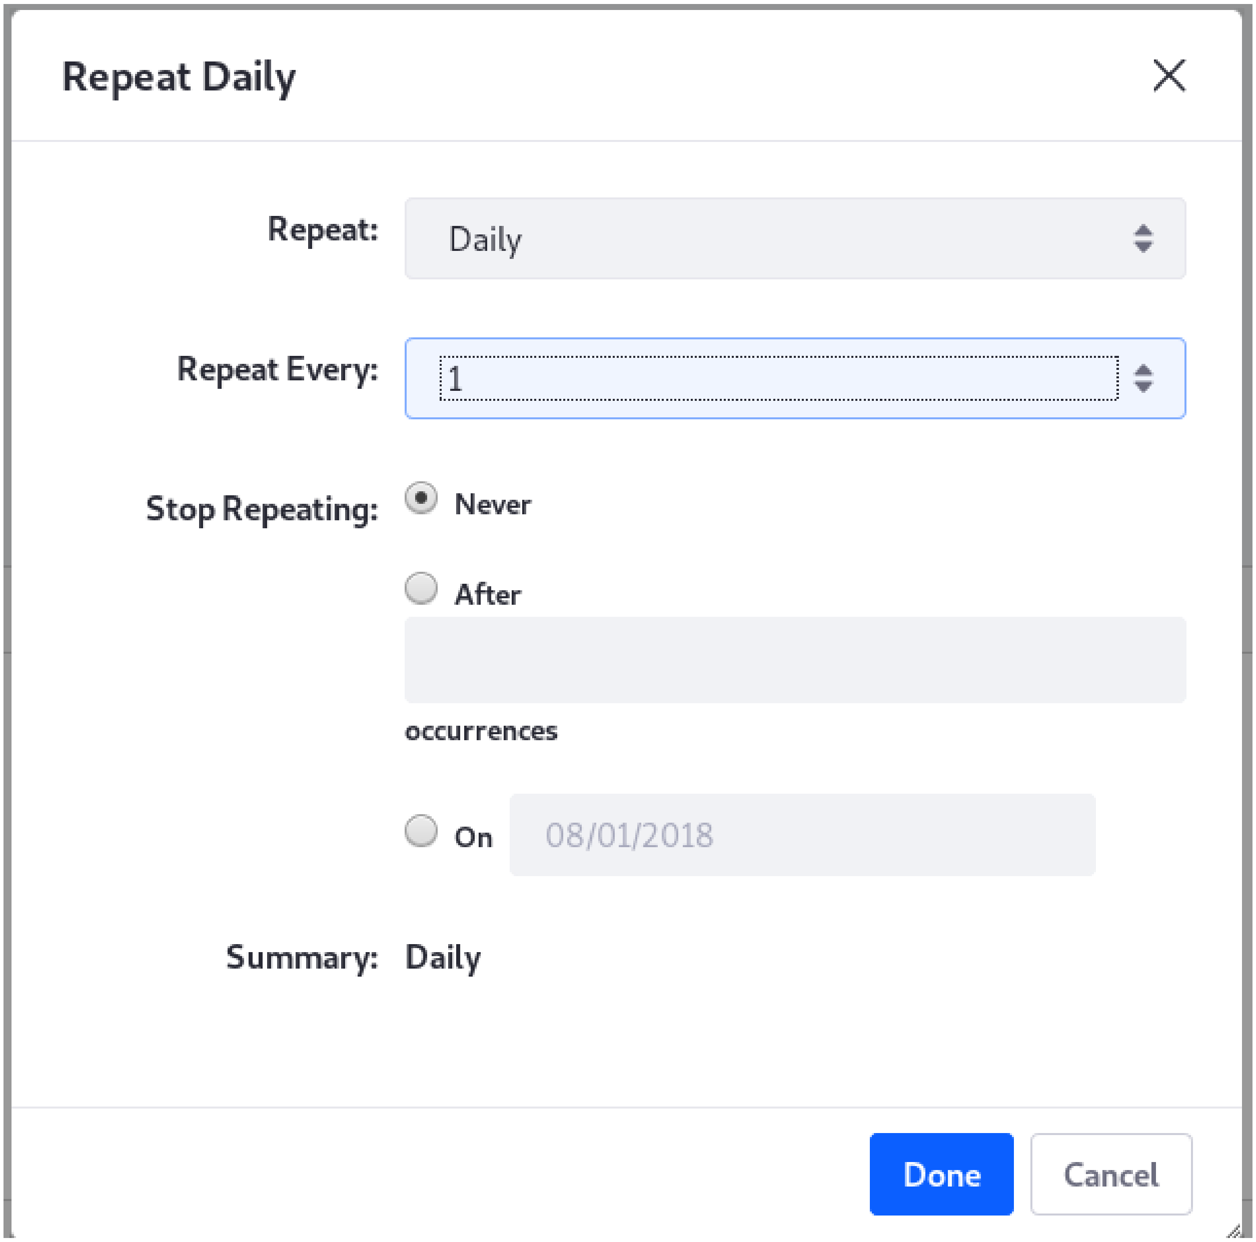

Figure 78.6: The *Repeat* box allows you to specify whether an events repeats daily, weekly, monthly, or yearly, how often it repeats, and when (or if) it ends.

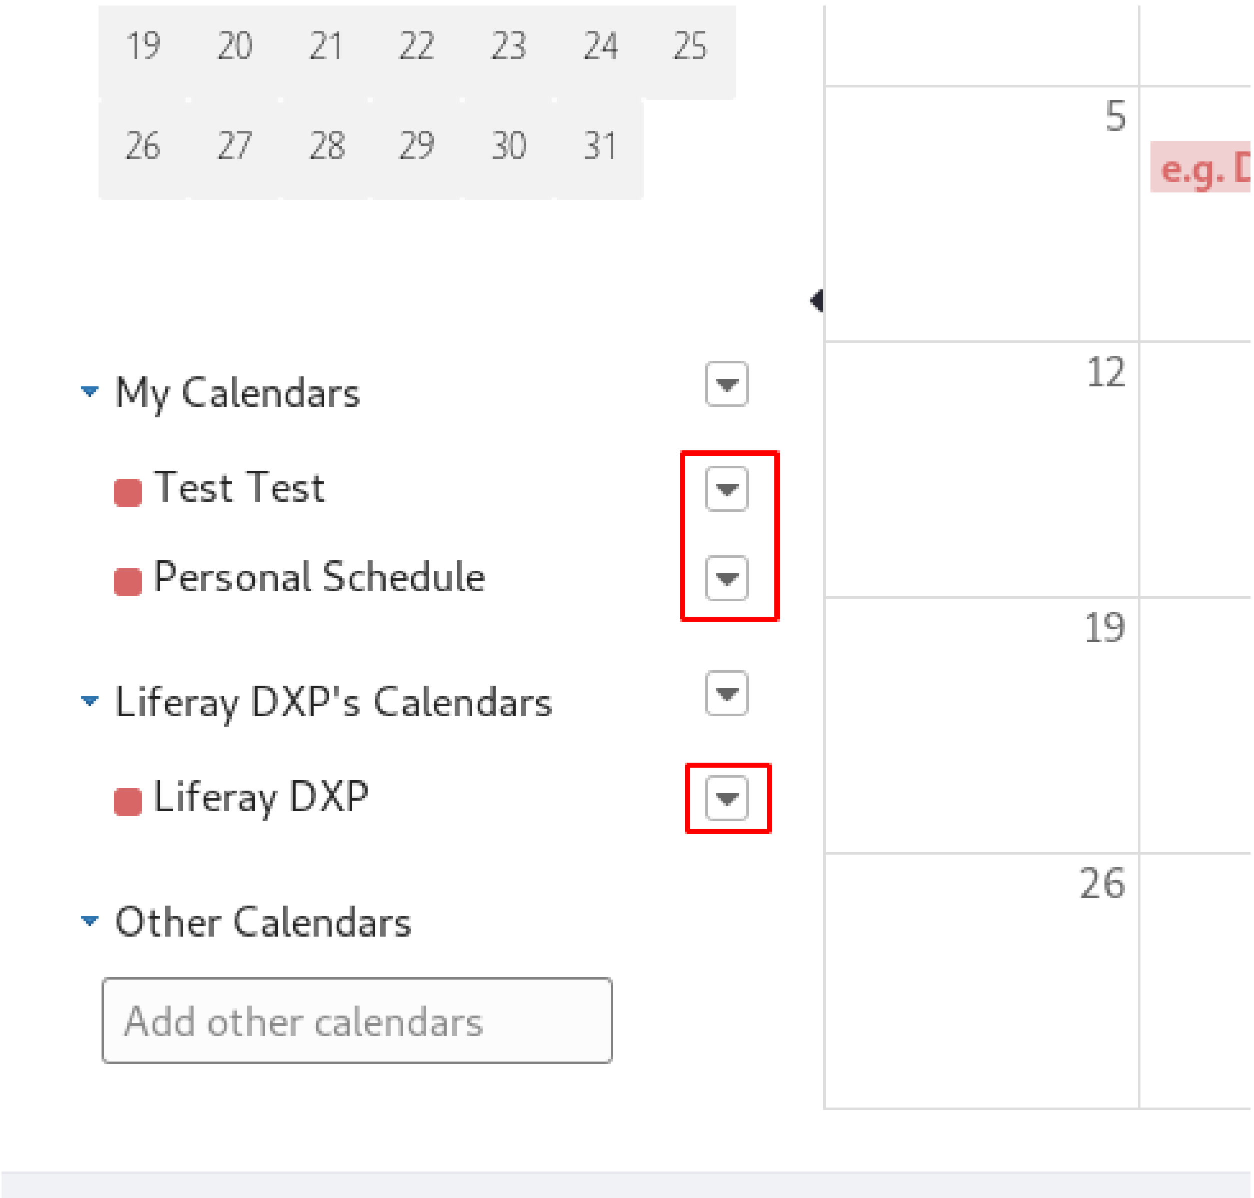

Figure 78.7: Email templates apply to a single calendar and all its events.

Calendar Resources

Follow these steps to add a new calendar resource:

1. Click on the *Resources* tab and click *Add Resource*.

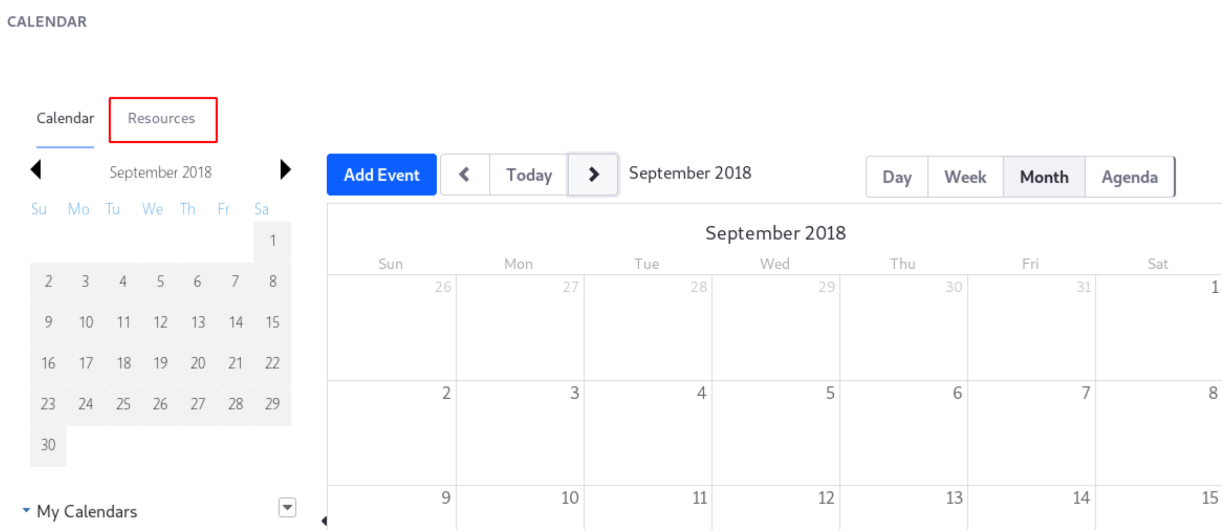

Figure 78.8: Resources are accessed from the tab menu at the top of the widget.

2. Fill in the *New Resource* form. Enter a name, give it a description, and choose whether to set it as active. You can also tag it, assign it to categories, and configure its permissions. Click *Save*.

The resource has its own calendar that was generated automatically (this is how users can check its availability when creating events). Just as with Users, however, resources can have more than one calendar. Follow these steps to assign a new calendar to the resource:

1. Go to the widget's *Resources* tab, click the ⋮ button next to the resource, and select *View Calendars*.

2. Click *Add Calendar* and continue just as if you were creating a calendar for a user or a Site.

Once a resource is created, invite it to your events just as you would an attendee.

Exporting and Importing Calendar Data

Like other Liferay Applications, the calendar allows data to be exported or imported as LAR files. As with all LAR files, data can only be ported between installations of the same version.

1. From the calendar widget, click the ⋮ button in the widget header and select *Export/Import*.

2. Enter a name for the LAR file (or use the default).

   Under *Application*, choose whether to include the widget's configuration in the LAR.

Under *Content*, choose how much historical data to export and select the content types (calendars, resources, and events) to include. You can also choose whether to include comments and ratings.

Check the appropriate boxes to select whether to include deletions and permissions in the LAR.

3. Click *Export*. When a success message displays (this may take a few moments) you can click on the LAR's filename to download it.

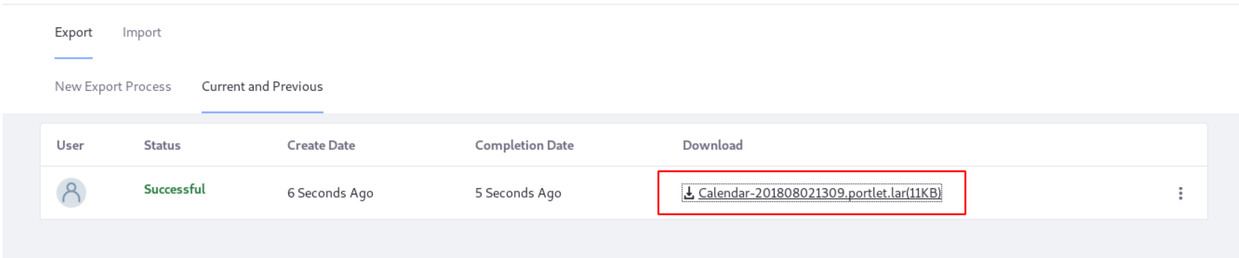

Figure 78.9: This LAR is ready to be downloaded.

Follow these steps to import a LAR:

1. From the calendar widget, click the  button in the widget header and select *Export/Import*.
2. Click the *Import* tab.
3. Click *Choose File* or else drag-and-drop a LAR into the area surrounded by a dotted line. Click *Continue*.
4. Decide how much data you want to import:

   Under *Application*, check the box to import the configuration stored in the LAR or leave in unchecked to keep your current configuration.

   Under *Content*, decide which content types (calendars, resources, and events) to import, and whether to include comments and ratings.

   Choose whether to import permissions and deletions, and decide whether to delete your widget's existing data before the import.
5. In the collapsible *Update Data* section, choose how data will be updated.

   Mirror: The data will be imported along with a reference to its source. This allows data to be updated rather than duplicated if the same LAR is imported more than once.

   Copy as New: All data is imported as new entries. Repeat imports will produce duplicates.
6. In the *Authorship of the Content* section, choose whether to keep the original author of the imported content (where available) or to list the current user as the author.
7. Click *Import*.

Your calendar is set up and ready to go! Better check it to see what's next on the agenda.

Part II

Discovering Analytics Cloud

CHAPTER 79

GENERATING NEW BUSINESS USING ANALYTICS

Your company interacts with lots of people, including existing and prospective customers. On your websites they browse products and services, ask questions, leave feedback, and more. Presenting relevant, interesting site content motivates people to do business with you. Knowing your customers well and understanding how they interact with your digital content can help you improve their experience.

Liferay Analytics Cloud (Analytics Cloud) is a cloud-based SaaS offering that works with Liferay Digital Experience Platform (Liferay DXP) to give you insights into customer attributes and behavior. It aggregates your data from multiple sources and presents analytics, making it easier to understand people and their interests. The centralized contact data and analytics help you communicate intelligently with people and create terrific user experiences that drive business and make people happier.

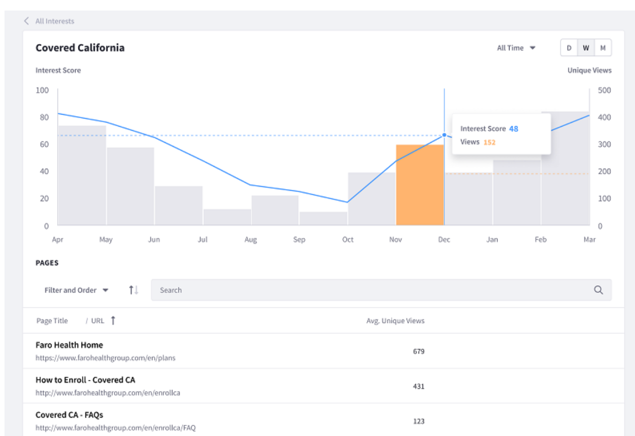

Figure 79.1: As you learn about topics that interest your audience, you can speak intelligently about the topics, through online content and in person.

79.1 Benefits

Here are key things Analytics Cloud helps you do:

- **Gain a well-rounded view of people:** Analytics Cloud gives you a clear picture of people by aggregating their profile properties and behavior data from multiple sources, including Liferay DXP sites.

- **Identify and track target audiences:** You can focus on characteristics and behaviors common to targeted individuals, aggregate those individuals into Segments, and track Segment growth.

- **Learn people's interests:** Once you know topics that interest people, you can engage them on those topics. You can also produce stimulating content that highlights how your products and services relate to the topics. These things help you generate new business and improve customer relationships.

- **Measure content performance:** Identify content that gets the most activity so you can build off its success by creating similar content to grow other parts of your business. Spot under-performing content or content issues that prevent people from engaging more deeply with your brand. Fix or improve the content to get people on track, or remove the content from your site.

- **Determine the best advertising channels:** Discover how customers are finding your site pages so you can advertise on the right channels.

These are just some ways you benefit from using Analytics Cloud. Before getting started with Analytics Cloud, you should understand its key components.

Analytics Cloud helps you understand people, their interaction with your sites, and the paths (or channels) that draw them to your sites in these ways:

- Understanding People

- Touchpoint Analytics

- Path Analytics

Start with understanding people.

79.2 Understanding People

Existing and prospective customers interact with your company in person, on your websites, and via email and messaging. On Liferay DXP sites, people are registered users or visitors. Analytics Cloud integrates contacts (people) from multiple sources, including Liferay DXP sites, into a single view. It merges contact data into rich profiles of *Individuals*—unique contacts. You can name contact fields your way, making them most useful to you and your company.

Analytics Cloud lets you aggregate Individuals into *Segments*. Segments are groups of Individuals with common characteristics. You can create Segments dynamically based on criteria such as industry and location, or statically (comprising hand-selected Individuals). Industry Segment Analytics shows industries engaging with you most and identifies your content that interests them.

Contact profiles are a big part of Analytics Cloud, but analytics about people's interaction with your website pages and assets is a bigger part. This type of analytics (called Touchpoint Analytics) helps you determine how effective your content is in achieving your business goals.

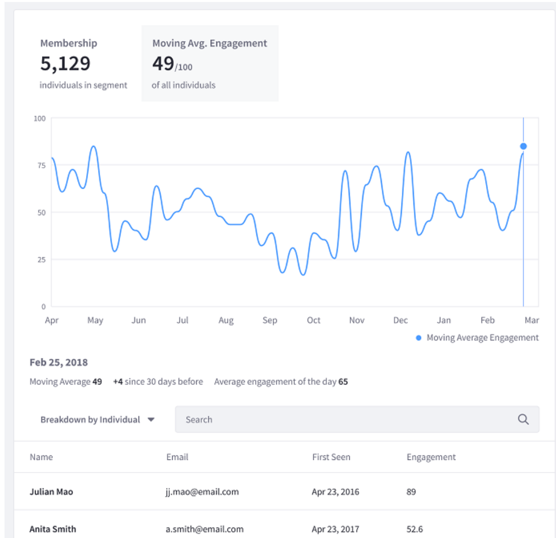

Figure 79.2: Dynamic Segment membership growth (and decline) reflects the number of people matching the Segment's criteria.

79.3 Touchpoint Analytics

Touchpoints are the locations (pages) where people interact with content. Here are some of the Touchpoint Analytics:

- *Visitor count:* the number of users (guest and signed in) visiting a Touchpoint.
- *Session length:* amount of time people are spending on a page's site.
- *Number of sessions:* number of different people on a page's site.
- *Time on page:* amount of time people are spending on a page.
- *Bounce rate:* percentage of visitors to a particular website who navigate away from the site after viewing only one page.

Assets are specific types of content entities on site pages. Traditional web analytics are based on static websites—everyone sees the same page content. By additionally providing analytics on content articles, Analytics Cloud helps you determine the article effectiveness. It also provides insight into people's interaction with personalized content. It generates charts on target audience engagement with the content.

Liferay Analytics Cloud (beta) supports these Liferay DXP Assets:

- Blogs
- Documents and Media
- Forms
- Web Content

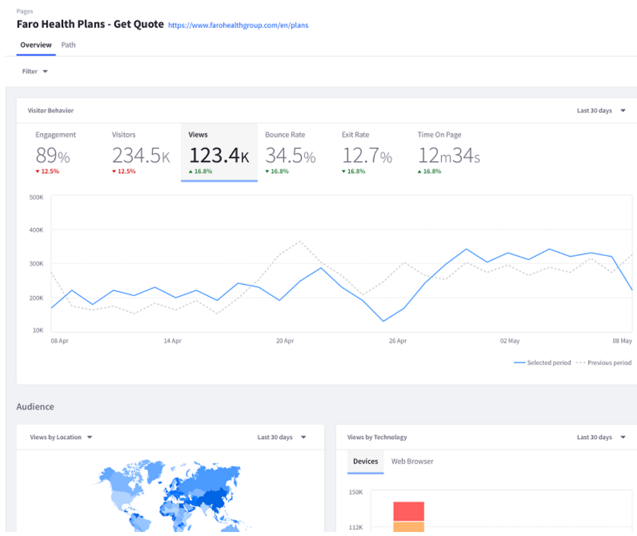

Figure 79.3: See how the number of views, visitors, and more are trending on your pages.

Regarding user behavior with a form, for example, Analytics Cloud shows the number of views and submissions, the percentage of users abandoning the form, and the form's average completion time. A line graph shows the number of views over time.

As you view Page and Asset Analytics you can discern whether an action is positive or negative. Here are some examples.

Positive:

- staying active in sessions
- frequent sign in
- submitting a form
- giving an article five stars

Negative:

- signing out shortly after signing in
- passing by key strategic content
- abandoning a form
- letting a session end

Learning how people get to pages and Assets is also useful. That's Path Analytics.

79.4 Path Analytics

People get to content through different channels (paths):

- Google search
- Advertisements
- Site navigation
- Search within the site

- Links inside other pages
- Links within Assets

If you expect your advertisements on a channel to draw lots of people to a page, for example, you can check its performance in the page's Path Analytics. You may learn that a particular social network increasingly drives site traffic, and thus advertise more on that network.

Path Analytics reports the device and browser types that render your site content. It helps you determine which devices and browsers to test more and optimize user experiences on.

You've been introduced to Liferay Analytics Cloud and what it offers: Segments, Touchpoints, Paths, and more. It's time to experience Analytics Cloud for yourself.

Chapter 80

Your New Home for Analytics

To get started using Liferay Analytics Cloud, the first thing you need is an invitation. Liferay invites the Analytics Cloud owner (the person who purchased the product) to your company's Analytics Cloud project. If you're not the owner but still need access to the project, ask an Analytics Cloud system administrator to invite you.

Your invitation links you to the Analytics Cloud sign in page, which prompts you for your Liferay credentials.

First time users see the Workspace creation page:

1. Name your workspace, select a server location for hosting your data, and review/agree to the terms and conditions.

2. Create the workspace. A message appears, stating that your Analytics Cloud environment is being set up. Next your project home page appears.

Subsequent visits to https://analytics.liferay.com/ show the Workspace welcome page, where you can select your workspace or any other accessible workspaces for working on projects.

The project home page describes and links to each key component page:

Segments: Aggregate customers into Segments based on similar profiles and behavior data. Track Segment growth and discover patterns of interest. Discover your target audience.

Accounts: Learn how individuals that belong to Salesforce accounts interact with your sites. By combining Salesforce account data with data from other sources, Analytics Cloud presents a holistic view of the account and those in it.

Individuals: Learn how individual customers interact with your Sites. View holistic profiles that combine customer's profile data with their page and asset interactions for deep insight into their interests.

Sites: Combine path analytics and traditional page analytics to understand how customers interact with your content. Measure individual web page performance to improve digital experiences.

Assets: Track and visualize engagement levels of digital assets such as forms, blogs, web content, documents, and media to ensure that assets are driving conversions.

Data Sources: Integrate your customer profile and behavior data from multiple Liferay DXP sites and custom data sources.

Additional shortcuts and overviews are displayed on the right:

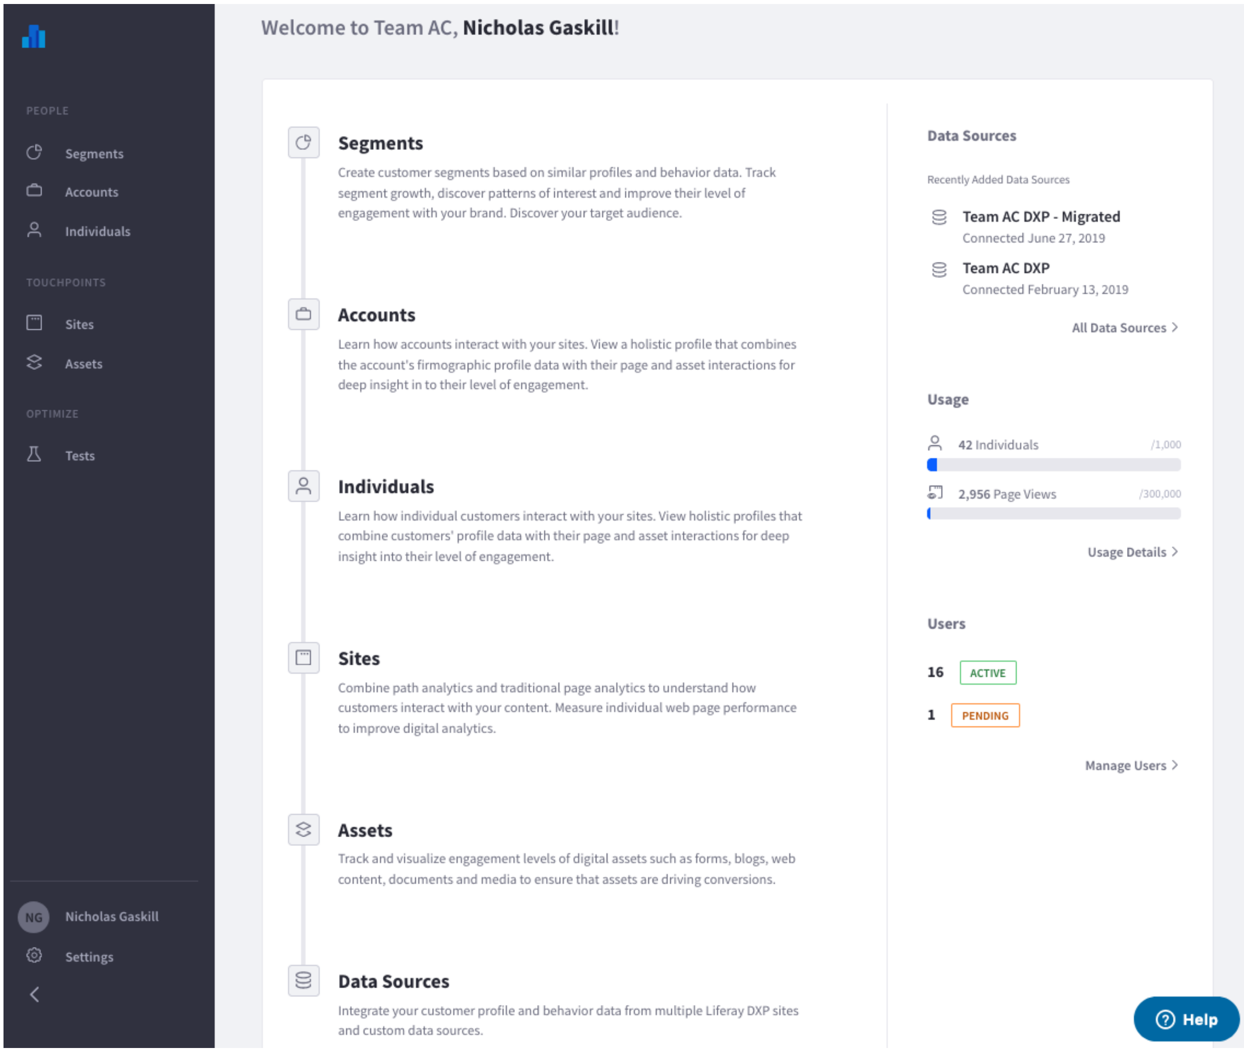

Figure 80.1: Your project's home page welcomes you to all Analytics Cloud offers.

Add Data Source: Quick access to adding data sources for customer profiles and behavior data.

Usage: Compares the number of Individuals and Pages you're analyzing to your Liferay Analytics Cloud Plan.

User Management: Invite users and assign permissions.

The *Liferay Analytics Cloud* icon ( ) at the top of the navigation panel takes you to the home page from wherever you are in the application. Beneath it are several other links:

People: Learn more about your customers by analyzing them individually as *Individuals*, aggregated as *Segments*, or as part of Salesforce *Accounts*.

Touchpoints: Analyze *page* and *asset* performance and view paths users take to your pages.

User icon: Switch projects or sign out.

Settings: Provides options for managing data sources, users, and updating your Analytics Cloud usage.

If no data sources are connected, a message near the top of the main area tells you so and links you to the page for adding data sources. Adding data sources is next.

80.1 Managing Data Sources

Liferay Analytics Cloud requires two kinds of data. First, you need web analytics data on user interactions with Liferay DXP Pages and Assets. Second, you need profile data on the users themselves. Analytics Cloud can sync these two kinds of data so that you can see not only how users are interacting with your site, but also who those users are. First, however, you must provide it with appropriate data sources.

Both kinds of data can be obtained from a Liferay DXP instance. If you have other user profile data—in addition to what is already stored in Liferay DXP user models—you can import it from a CSV file. In a future Analytics Cloud release, Liferay plans to support syncing with contact data from Salesforce too.

The contact data is merged and consolidated into a single customer view. Liferay DXP data sources let you select Organizations and User Groups to sync. Individual contacts are matched by email address. You can define the contact data model using whatever field names and types you want, but Analytics Cloud makes it easy to define trivial model fields by suggesting values.

Your data sources can be viewed from the navigation panel.

1. Select *Settings*.
2. Click on *Data Sources*.

The Data Sources page appears and lists all existing data sources.

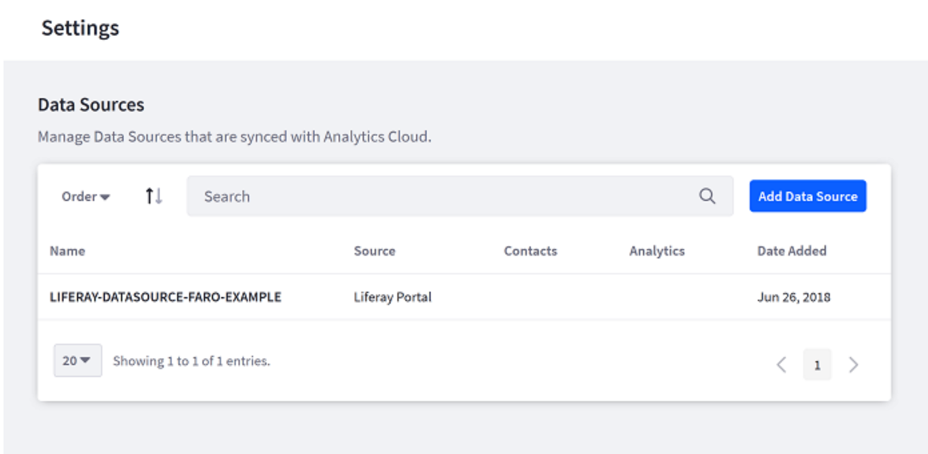

Figure 80.2: View, edit, and add data sources from the Data Sources page.

Unless a teammate has already added a data source, the list is empty. To add a new data source, see the following tutorials:

- Adding a Liferay DXP Data Source
- Adding a CSV Data Source

Once you've created your data sources, you might need to modify them from time to time.

Modifying a Data Source

The Data Sources lets you page through the data source listing or use *Search* to find one that matches a keyword. Once you've found the data source, click on it to edit it. Here's what you can modify:

- Name of the data source
- Contact mapping
- Enable/disable Contact or Analytics sync

Here's how to delete a data source:

1. Click on the listed data source. The data source's page appears.

2. Click *Delete Data Source*.

Your Liferay DXP instances are rich sources of customer profile and behavior data. Adding a Liferay DXP instance as a data source is next.

80.2 Adding a Liferay DXP Data Source

Your Liferay DXP instances are rich with contact data from Users and web analytics data on user interaction with Liferay DXP pages and assets. When you Create Liferay DXP data sources, you can select web analytics data from the Liferay DXP Sites you want. In Users and Organizations, you can choose to use contact data from all Users or a specified subset of them. Before adding an instance as a data source, though, you must connect it to your Analytics Cloud project.

Liferay DXP Data Source Prerequisites

Follow these steps to connect your project to your Liferay DXP instance.

Step 1: Install Required Liferay DXP Fix Packs

Liferay DXP 7.1: Download and install fix pack 10+.
Liferay DXP 7.0: Download and install fix pack 79+.

Step 2: Register Analytics Cloud with your Liferay DXP instance

Liferay DXP 7.1:

1. Download the Liferay Plugin for OAuth 2.0 version 1.1.0 (or newer) and install it.

2. The plugin comes with Analytics Cloud pre-registered. Copy the *Client ID* and *Client Secret* for connecting DXP with Analytics Cloud, as described in the next section.

Liferay DXP 7.0:

1. Download the Liferay Connector to OAuth 1.0a and install it.

2. Register Analytics Cloud as an OAuth application with the *Write* access level.

3. Copy the *Consumer ID* and *Consumer Secret* for connecting DXP with Analytics Cloud, as described in the next section.

Congratulations on authorizing Analytics Cloud to connect to your Liferay DXP instance! It's time to add your DXP instance as a data source.

Tip: If you have problems connecting to your DXP data source, refer to Troubleshooting Liferay DXP Data Sources.

Adding the DXP Data Source

Adding a Liferay DXP data source connects your Analytics Cloud project with a Liferay DXP instance.

1. Select *Settings → Data Sources*. A listing of your data sources appears.

2. Click *Add Data Source*. The *Add Data Source* page appears.

3. Select the *Liferay DXP* icon. The *Configure Liferay DXP* page appears.

The Authorization tab is selected by default. It's time to authorize your DXP instance as a data source.

DXP Data Source Authorization

Here's how to authorize your DXP instance as a data source:

1. Fill in the data source and client credentials fields.

   Description

   Name: A name for your data source.

   URL: The Liferay DXP instance URL.

   Client Credentials

   Consumer Key/Client ID: Key/ID for Analytics Cloud to access your Liferay DXP instance.

   Consumer Secret/Client Secret: Secret for Analytics Cloud to access your Liferay DXP instance.

   In Liferay DXP 7.1, the Client ID and Secret are found at *Control Panel → Configuration → OAuth 2 Admin*.

   In Liferay DXP 7.0, the Consumer Key and Secret are found at *Control Panel → OAuth Admin*.

2. Click *Authorize*. A window appears and prompts you to sign in to the DXP instance.

3. Sign in by entering your DXP admin (user that has the Admin role) credentials and clicking *Authorize*.

4. Click *Save* to save the authorization options. Analytics Cloud advances you to the Configure Data Source tab's Data Configuration page. The data source's Current Status is *AUTHENTICATED*.

Here are the Configure Data Source options:
Configure Contacts: Configures the contact data only.
Configure Analytics: Configures the assets and touchpoints only.
Start with configuring contacts.

Configuring Contacts

Configuring contacts imports DXP user data.

1. Select the *Configure* button for *Configure Contacts*. The Contacts configuration options appear.

2. Configuring contacts involves selecting contacts to sync from the Liferay DXP instance and its User Groups and Organizations. Contacts belonging to multiple User Groups and Organizations are only counted once.

   Sync All Contacts: Selects all Liferay DXP instance contacts and disables options for selecting specific User Groups and Organizations.

   Sync By User Groups: Selects contacts by User Group.

   Sync By Organizations: Selects contacts by Organization.

3. Click *Save and Continue* to import the selected contacts. Analytics Cloud imports the contact data and maps it to your Analytics Cloud contact data model. The initial contact data import can take 5 1/2 minutes per 1,000 contacts.

4. Follow instructions for Mapping Contact Data to map contact data from your Liferay DXP instance to your Analytics Cloud contact data model. Once you've mapped the data, click *Save*. The Data Configuration page appears again; the button for Configure Contacts is labeled *Edit*.

You've configured Analytics Cloud to use your Liferay DXP contacts.

Configuring Analytics

Configuring analytics imports asset and touchpoint data as it relates to DXP contacts you've imported.

1. Click *Configure* for *Configure Analytics*. The Liferay DXP site analytics registration page appears.

2. Select the Liferay DXP Sites to register for analytics and click *Configure*.

3. Click the *Done* button.

The Contacts and Analytics data start syncing into Analytics Cloud. **Initially the sync takes a while. After the initial sync, changes are synced periodically.**

Tip: If you have problems connecting to your DXP data source, refer to Troubleshooting Liferay DXP Data Sources.

If you have contact profile data from other sources such as a database, you might be able to export the data to a CSV file. Then you can add the CSV file as a data source and import the contact profiles from it. Adding a CSV data source is next.

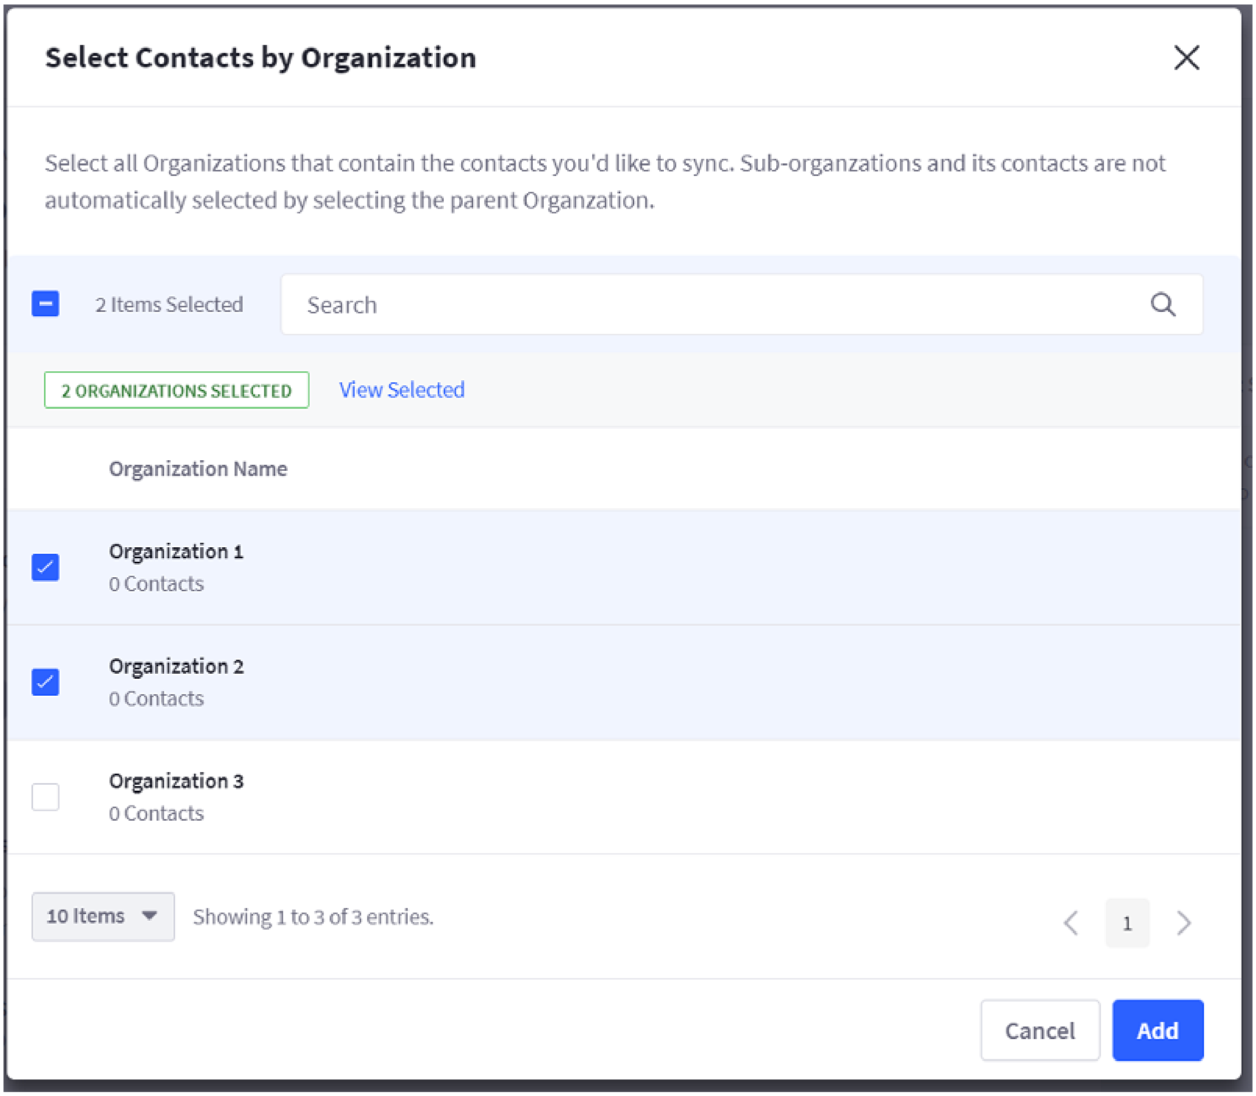

Figure 80.3: Analytics Cloud lets you select and import contacts from a Liferay DXP instance and its Organizations and User Groups.

80.3 Adding a CSV Data Source

You can import contact profile data from CSV files to enrich customer profiles with additional data related to a user's industry, job title, annual income, or whatever metrics are important to your business. If you have customer profile data in databases or collect it through web forms, you can export the data into CSV files.

Important: The CSV files must have an email column.

Here's how to integrate contact data from a CSV file:

1. In the Data Sources page, click *Add Data Source*. A page appears showing the data source types.
2. Select the *CSV File* icon. The CSV file upload page appears.
3. Upload your CSV file by dragging it into the upload area or browsing and selecting it from your file system. Click *Next* to upload the file. The CSV configuration page appears.
4. Configure details and metadata about your CSV file data source, and click *Next* when you're done.
    1. *Name:* Enter a name for your data source.
    2. *Context:* Select a context that describes where the contact data is from. On selecting context, fields appear for you to further describe the data.
    3. *View Data Preview:* Shows the raw data in table format. If it's not formatted the way you want, adjust your CSV file's format.
5. Follow instructions for Mapping Contact Data to map contact data from your CSV file to your Analytics Cloud contact data model. Once you've mapped the data, click *Next*.

The contact profile data starts syncing. The time it takes to sync depends on the number of contacts.

From time to time you might need to modify your data source configurations. If you've already explored *Mapping Contact Data* (the next topic) skip ahead to *Editing Data Sources*, to see how easy it is.

80.4 Adding a Salesforce Data Source

Analytics Cloud can integrate Salesforce data with data from Liferay DXP and other sources, automatically blending the data to build a comprehensive customer profile. This lets you create more accurate segmentation and enhance your site's personalization strategy. To do this, you must first add a Salesforce instance as a data source.

Adding the Data Source

1. Select *Settings* → *Data Sources* → *Add Data Source*. This opens the Add Data Source screen.
2. Select *Salesforce*. This opens the Configure Salesforce screen.

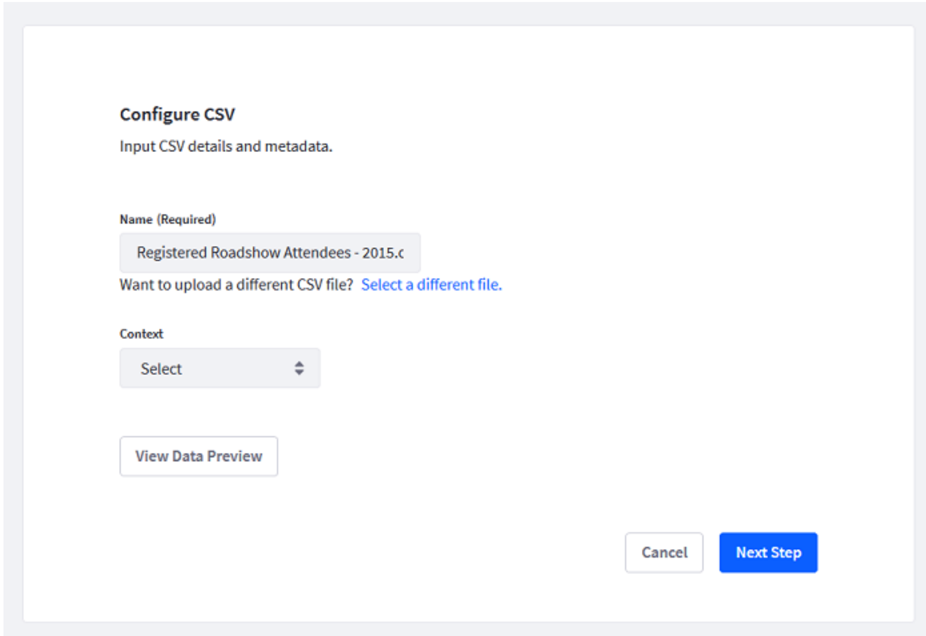

Figure 80.4: When configuring a CSV file data source, you can describe the data context and view the data to make sure it's formatted properly.

3. In the Authorization tab, enter the Salesforce instance's name and URL. Under the CLIENT CREDENTIALS section, enter the client ID and client secret of the OAuth connected app in Salesforce. Note that the Salesforce instance's administrator must create this connected app with the following settings:

   Callback URL: `https://analytics.liferay.com/oauth/receive`

   OAuth Scopes:

   - Access your basic information (id, profile, email address, phone)
   - Access and manage your data (api)
   - Perform requests on your behalf at any time (refresh_token, offline_access)

   For instructions on creating an OAuth connected app in Salesforce and locating its client ID and client secret, see the Salesforce documentation.

4. Click *Authorize & Save*. This begins importing the leads, contacts, and accounts data from Salesforce. This data is integrated with Analytics Cloud data as follows:

   - Data from Salesforce leads and contacts are merged with data for matching individuals in Analytics Cloud. The match is based on email address. For example, if a contact and/or lead has the same email address as an individual in Analytics Cloud, the data of that contact and/or lead is merged into that individual's Analytics Cloud data.
   - If the email address of a lead and/or contact doesn't match with that of an existing individual in Analytics Cloud, a new individual is created in Analytics Cloud with the data from the lead and/or contact.
   - Data from a contact is prioritized over that of a lead. For example, if the same field is populated in a matching lead and contact, then the data from the contact is imported into the individual in Analytics Cloud.

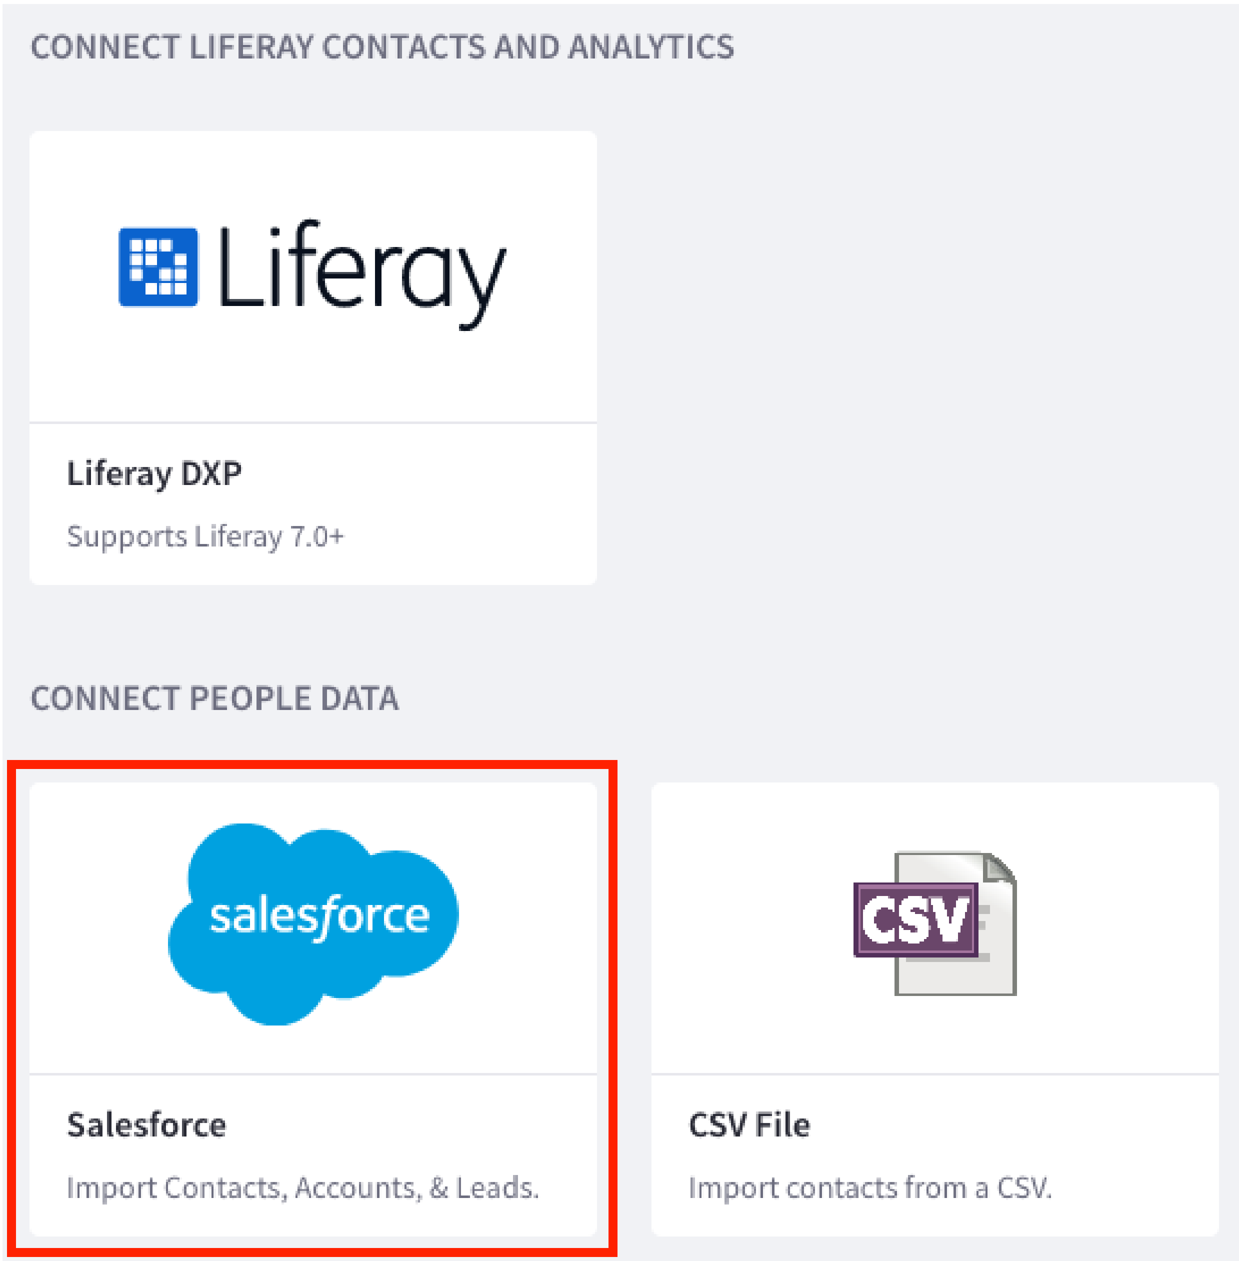

Figure 80.5: Select Salesforce from the Add Data Source screen.

Note that importing the data may take some time, depending how much data exists in the Salesforce instance.

5. The import progress is visible on the *CONFIGURE DATA SOURCE* tab. Click *View* to see the mapping of each Salesforce field (for Accounts and Individuals) to its corresponding Analytics Cloud field. You can also view this information later from *Settings* → *Data Sources* → *(Your Salesforce)* → *CONFIGURE DATA SOURCE*.

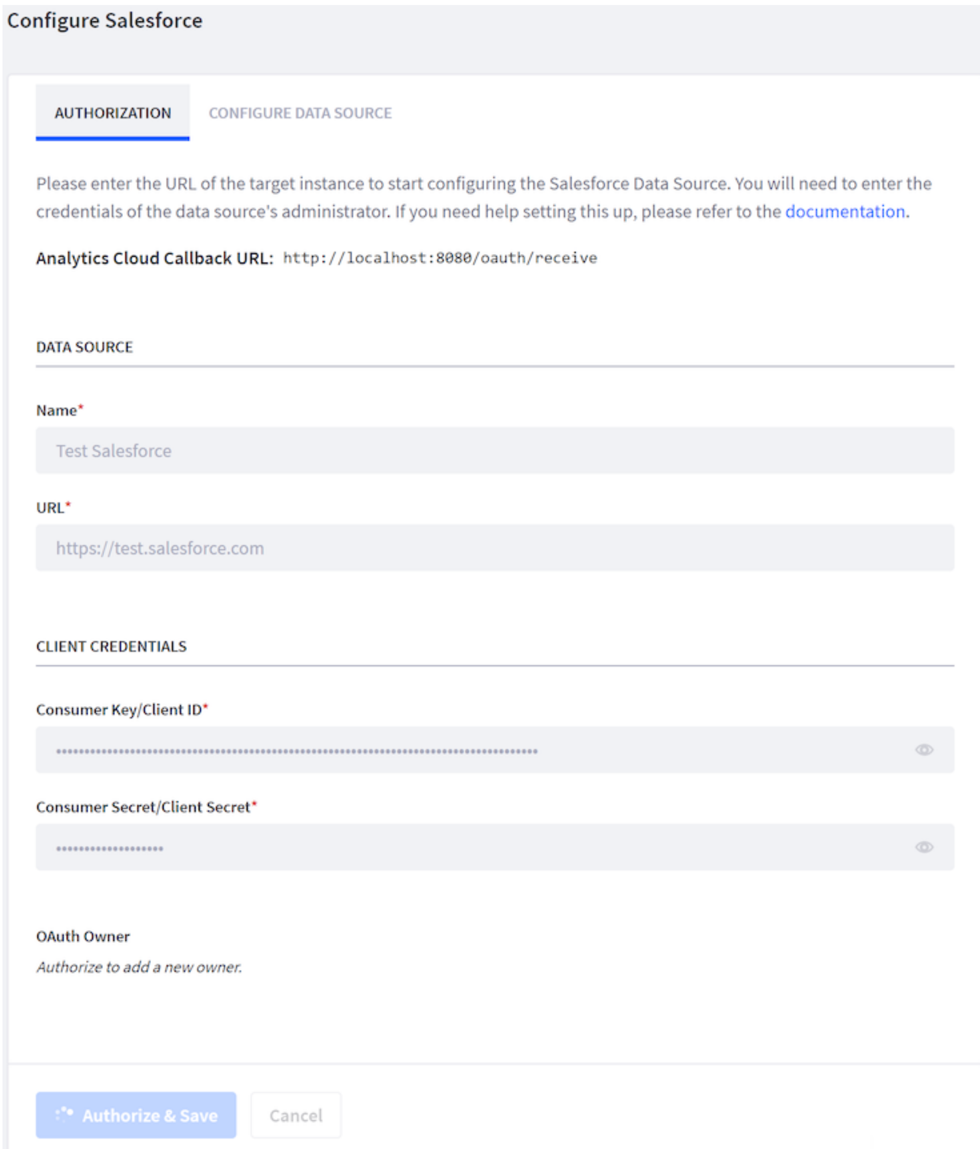

Figure 80.6: Enter the information needed to connect to your Salesforce instance.

80.5 Mapping Contact Data

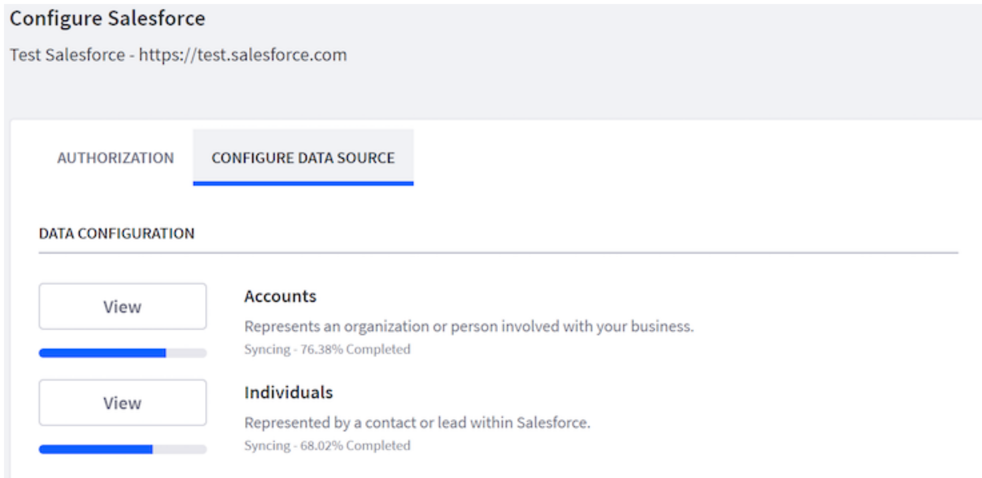

Figure 80.7: The CONFIGURE DATA SOURCE tab shows the status of the accounts and individuals imported from Salesforce, as well as the field mapping.

Analytics Cloud can map contact profile fields into a unified customer data model. It starts you with a default model. The process for mapping contact data is similar across data source types. Once you've selected the contacts to sync, Analytics Cloud makes a best effort to map contact data fields from the data source (e.g., Liferay DXP or CSV) to your unified contact data model.

Analytics Cloud provides several contact data mapping options:

- Choose the most appropriate data model field—Analytics Cloud *suggests* possible matches.
- Add a new custom data model field and map source data to it.
- Map source fields from multiple data sources to the same data model field.

Note: If you map source fields from multiple data sources to the same model field, the latest modified value is used.

As you map source fields to the data model, you can search for and select from the data model fields, use a *suggested* field, or create a new custom field.

Here's how to create custom contact data model fields:

1. Click on the data model field's selector.
2. Click *New Field*. A dialog appears for you to create the new field.
3. In the dialog, name the new field and select its type.
4. Click *Create*.

Your custom model field is ready to match with source fields.

When you're done mapping the data, click the *Done* button. Congratulations on mapping contacts from your data source to your Analytics Field contact model.

Once you've mapped your contact data and finished adding your data sources, you can start learning more about your customers. See Understanding People for more.

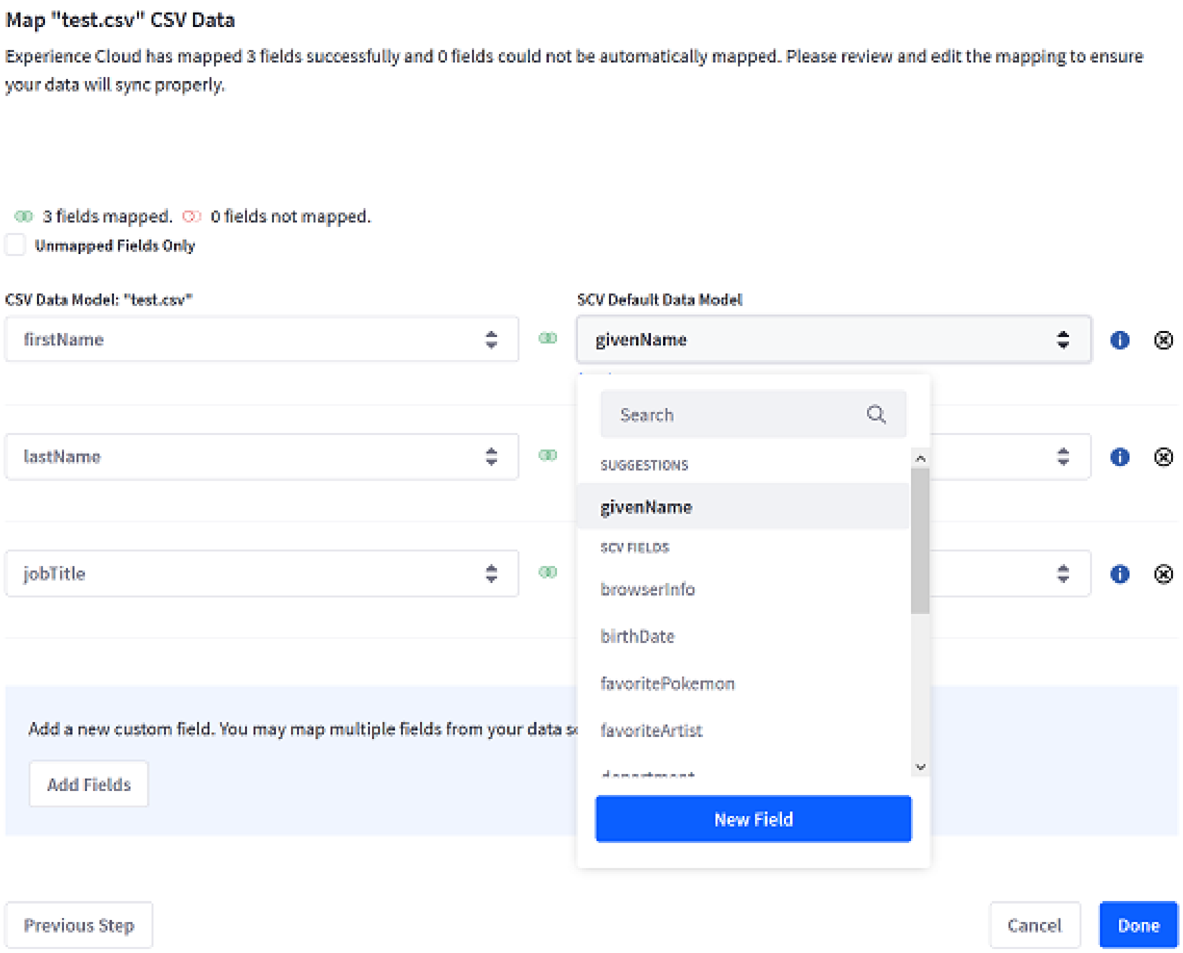

Figure 80.8: Analytics Cloud facilitates finding appropriate data model fields and offering suggestions.

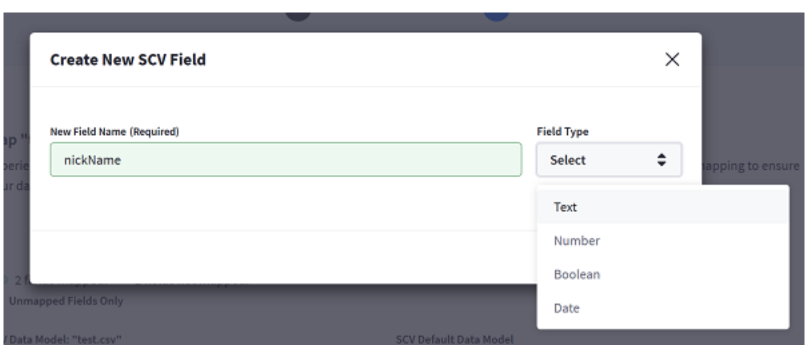

Figure 80.9: Creating new data model fields is easy.

Chapter 81

Learning What Makes Customers Happier

Analytics Cloud creates a single source of truth that brings together the various facets of customer knowledge to paint a more rounded, more informed picture of each customer and his or her behavior. It gives you a single customer view.

Here, your business users access a single location in which to view and manage customer data compiled from your customer data sources—DXP sites and customer CSV files. Analytics Cloud can merge customer data your way, into a model that works best for you.

This foundation reveals insights into individual customer behavior, such as how often individuals visit your sites, the content they read, and the products they browse. The analytics identifies positive and negative customer experiences. While you and your teammates might individually know parts of these things, providing this information in one place allows for highly targeted communication and marketing efforts.

81.1 Profiling Individuals

Many individuals visit your sites. Learning who they are and their activities and interests helps you improve business with them. Individual profiles help your Sales, Customer Service, and Customer Relations teams work personally with them. The profiles help you gain more context on what people do on your websites and their interests, so you can communicate with them intelligently according to their needs.

Here's how to view Individual profiles:

1. Select *Individuals* from the left navigation panel. The *Individuals* page appears, listing individual contacts, so you can page through, search for, and select them to learn more about them.

2. Select an Individual to view these things about the person:

    - Activities on your sites
    - Identity details and properties
    - Topic interests
    - Segment membership

Each Individual profile has these tabs:

- Overview
- Details
- Interests
- Segments

Let's explore them to learn more about people you do business with.

Overview

The Overview tab shows the person's recent activities using a histogram of the customer's activities and frequency in the past 90 days. The activities include page views, downloads, and form submissions. If you click on a specific day, it lists the Touchpoints (pages and assets) the customer acted on during that day only.

The Overview tab also gives you a glimpse of a person's profile data, topics of interest, and associated audience segments. Clicking the *View all ...* link for the Details, Interests, or Segments summary opens its respective tab.

Click on the *Details* tab to explore an Individual's details.

Details

The Details tab shows all of the person's attributes, the data sources they came from, and when they were last modified.

Interests

Click on the *Interests* tab.

Depending on the web pages or assets a user interacts with, Analytics Cloud uses an advanced algorithm to find out what an individual's interests are with your brand. This helps you prepare for conversations with the individual or create campaigns based on the interests.

See Customer Insights for details on the Interests insight.

Segments

The Segments tab lists all the Segments the customer belongs to. Analyzing Segments helps you learn about the Individual and similar Individuals. Satisfying Segments of customers generates new business fast. Segments are explained in detail next.

81.2 Creating Segments

Segments aggregate Individuals based on common properties and behavior. For example, you can create a Segment that contains users who are interested in a given topic, work in a specific industry, or both. Analytics Cloud then analyzes and profiles Segments. You can then integrate your Analytics Cloud Segments with personalization in Liferay DXP. This lets you deliver content of interest to each Segment via Liferay DXP.

Here, you'll learn how to create segments in Analytics Cloud. In Analytics Cloud's navigation panel, *Segments* is at the top of the *People* section because it's where you'll spend most of your time gaining customer insights.

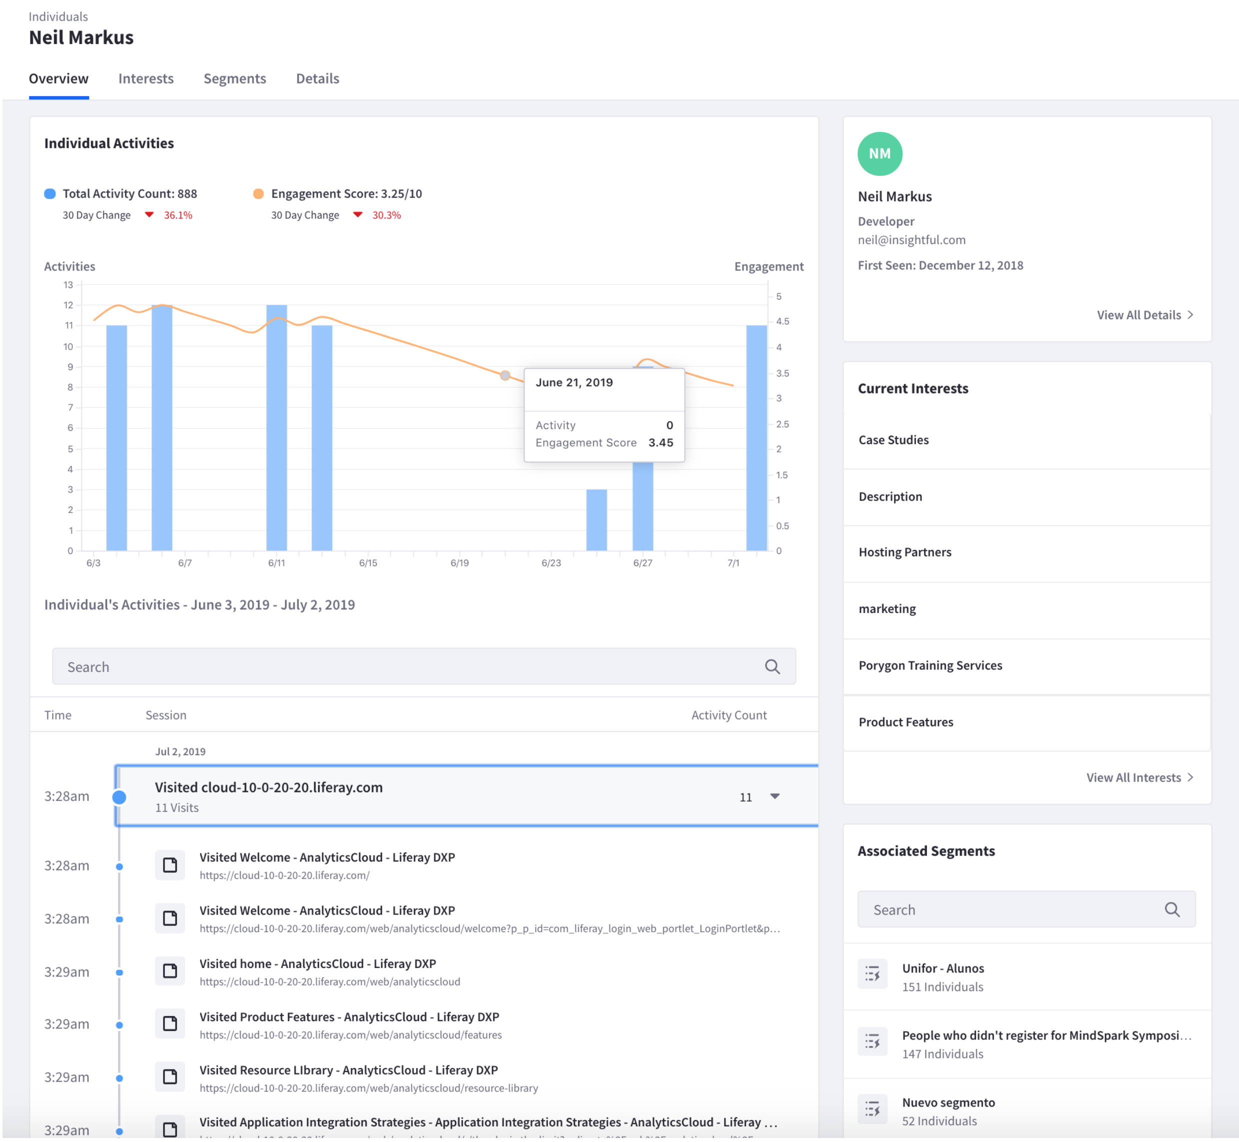

Figure 81.1: An Individual's Overview tab is a great place to learn how a person is interacting with your business.

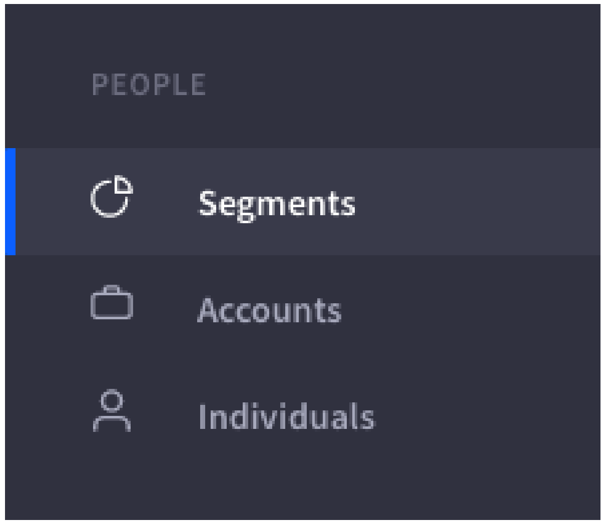

Figure 81.2: You can access Segments in the navigation panel.

There are two types of Segments:

Dynamic Segments: Individuals aggregated automatically based on criteria that you specify. Criteria can be based on Individuals' properties and interaction with your Liferay DXP Site pages and assets. For example, you can use an *industry* property as criteria. Analytics on industry Segments show which industries engage with you the most, and what content they're interested in.

Static Segments: Individuals aggregated manually. Static Segments are therefore comprised of hand-selected individuals.

Creating Dynamic Segments

Follow these steps to create a Dynamic Segment:

1. Click *Segments* in the navigation panel. A table with existing Segments appears.

2. Click *Create Segment* → *Dynamic Segment*. The *Create Individuals Segment* screen appears for creating a Dynamic Segment.

3. Click the *Edit* icon next to *Unnamed Segment*, and give your segment a name.

4. To create your Segment's criteria, drag and drop them from the panel on the right to the canvas on the center of the screen. The selector menu at the top of the panel lets you select from these criteria types:

   Individual Properties: Properties that belong to a person (e.g., age, country, industry, etc.).

   Account Properties: Properties that belong to a Salesforce Account (e.g., account name, industry, website, etc.).

   Interests: Topics and content types that interest the person. For more information on interests in Analytics Cloud, see Customer Insights.

   Session Properties: Properties that belong to a person's web session (e.g., browser, geolocation, etc.).

   Web Behaviors: Actions taken by a person (e.g., submitted a form, viewed a blog, etc.).

For further instructions, see Creating Criteria.

5. Anonymous users are excluded from Segments by default. To include them, enable the *Include Anonymous* toggle. Note, however, that criteria with individual and/or account properties will exclude anonymous users regardless of your setting here. Such properties only apply to known users.

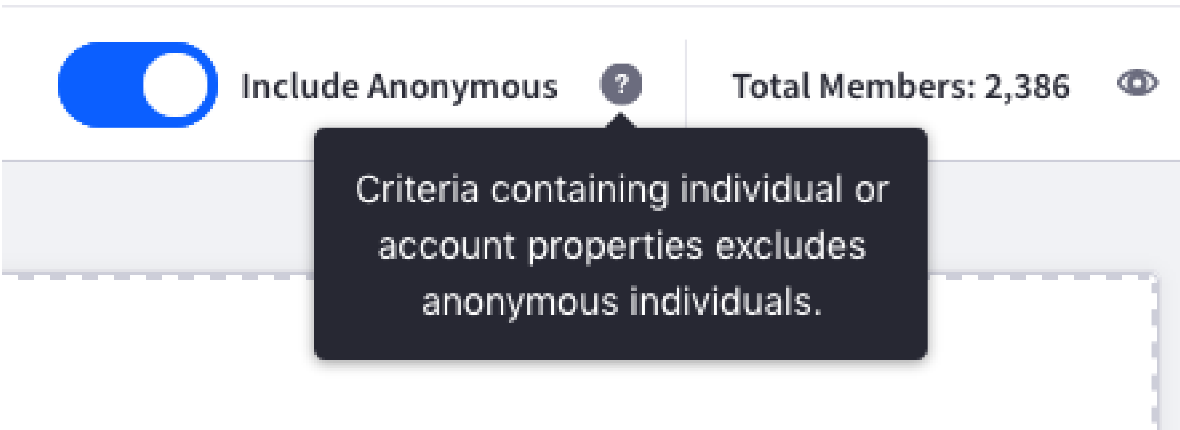

Figure 81.3: You can also include anonymous users in your segment.

6. Click *Save Segment* when you're finished.

Creating Criteria

The criteria creation canvas is very flexible. Once added to the canvas, you can move, delete, or duplicate any criterion:

To move: Click and drag the criterion using the vertical dots ( ) on its left.

To delete: Click the criterion's trash icon ( ). Alternatively, you can click the criterion's Actions icon ( ) and select *Delete*.

To duplicate: Click the criterion's Actions icon ( ) and select *Duplicate*.

Each criterion that you add contains fields that let you customize it to your needs. The first field is typically a selector menu in which you specify a condition for any remaining fields. The condition's values depend on the data type for the remaining fields. Here are some common condition values:

- Contains (text)
- Equals
- Greater than (number)
- Is known
- Less than (number)
- Does not contain (text)
- Does not equal
- Is not known
- Greater than or equals (number)

- Less than or equals (number)
- Is
- Is not
- Is before (date)
- Is after (date)
- Has (behavior)
- Has not (behavior)

For example, the `birthDate` criterion's first field is a selector menu that contains the options *is before*, *is*, and *is after*. The second field is a date field. You can therefore, for example, specify a criterion in which only Individuals with a birthday after 31 December 1980 are part of the Segment.

Figure 81.4: This criterion specifies Individuals with a birthday after 31 December 1980.

You can also control the way adjacent criteria interact with each other. For example, if you place criteria next to each other, a small box appears between them with the text *AND*. This means that the two criteria are joined by a logical *And*. Clicking the box changes it to *OR*, which represents a logical *Or*. Selecting *And* narrows the Segment's selection of Individuals; *Or* broadens it.

For example, joining two `birthDate` criteria with the following conditions creates a Segment targeting the Millennial generation (born 1981 - 1996):

- Is after 31 December 1980
- AND
- Is before 01 January 1997

You can also form subgroups of criteria by dragging and dropping criteria onto each other. An AND/OR box then appears between the subgroup and any adjacent criteria. Together, these tools let you build complex criteria for your Segment.

Creating Static Segments

Follow these steps to create a Dynamic Segment:

1. Click *Segments* in the navigation panel. A table with existing Segments appears.

2. Click *Create Segment* → *Static Segment*. The *Create Individuals Segment* screen appears for creating a Static Segment.

3. Name the Segment.

4. Click *Add Members* to bring up the Add Members screen, which contains a searchable list of all Individuals.

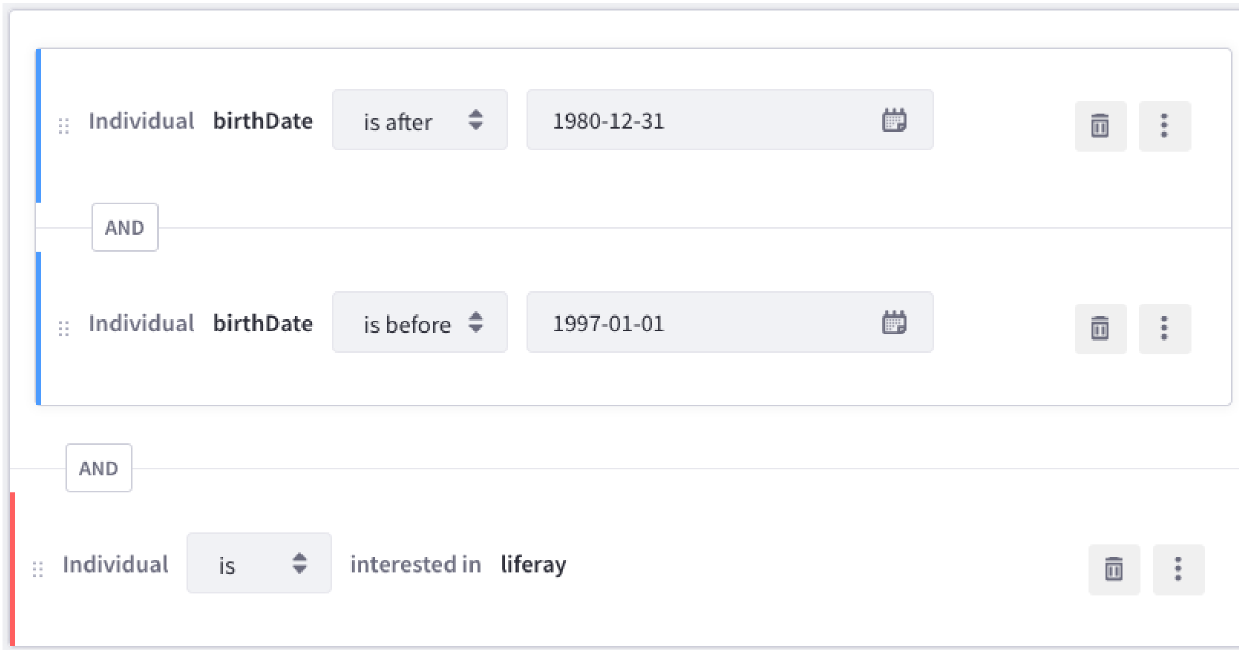

Figure 81.5: Millennials interested in Liferay love avocado toast.

5. Select the Individuals to add to the Segment, then click *Add*.
6. To change or undo your selections, click the *View Added Members* link and click *Undo* for each Individual you want to remove. Alternatively, select each Individual and click the *Undo Changes* button that appears. To remove all Individuals, click *Undo All*.
7. Click *Create* when you're finished.

81.3 Liferay DXP Personalization

When used with Liferay DXP's personalization features, the full power of Analytics Cloud's Segments is realized. Personalization lets you target Liferay DXP content at specific Segments. For example, if you have a Segment for users in the finance industry, then you could use personalization to show them content on investing.

Note that like Analytics Cloud, Liferay DXP can contain Segments. Liferay DXP segments, however, are less powerful than those in Analytics Cloud. Analytics Cloud Segments are more powerful because of the comparatively greater computing power and resources that exist in the cloud.

Here's a comparison of Segments in Liferay DXP and Analytics Cloud.

Liferay DXP Segments	Analytics Cloud Segments
Based only on user data in DXP	Based on user data from multiple sources (e.g., DXP, Salesforce, CSV files, etc.)

Liferay DXP Segments	Analytics Cloud Segments
Created from session attribute and cookies	Created from user interests and historical behavior
User identity comes from only one DXP site	User identity is resolved from multiple DXP sites
Can't include anonymous users	Includes anonymous users
Good for limited, short-term personalization	Good for expansive, long-term personalization

For information on using Analytics Cloud Segments in Liferay DXP, see Using Analytics Cloud with User Segments. For information on personalization in Liferay DXP, see Personalizing Experiences.

81.4 Profiling Segments

The Segment profiles are your gateway into how people interact with your sites. They show who's matching your Segment criteria (Dynamic Segments), so you can deliver content that wows them.

Here's how to view Segment profiles:

1. Select *Segments* from the navigation panel. The *Segments* page appears, listing individual Segments. You can page through, search for, and select Segments to learn more about them.

2. Select a Segment to learn these things about it:

    - *Membership:* Population fluctuation over time.

    - *Topics of interest:* Topics mentioned in the pages and assets the Segment visits most.
    - *Distribution:* Demographics based on Individual properties

The Segment's profile appears, showing its *Overview* panel.

Overview

The *Overview* tab's main area shows membership growth. As a marketer, you might want to generate more customer prospects in a specific industry. If you create a Dynamic Segment with that industry as criterion, you can view its membership chart to monitor customer growth in that industry.

The Overview tab also gives you a glimpse of the Segment's criteria, interests, and demographics. Each one links to more details, or you can click on their respective tabs.

Interests

As a Segment's members interact with your site pages and assets, Analytics Cloud's interests algorithm determines the level of interest the Segment has with your site's topics. Knowing these topics helps you prepare for conversing with the Segment members and create content that interests them.

See Customer Insights for details on the Interests insight.

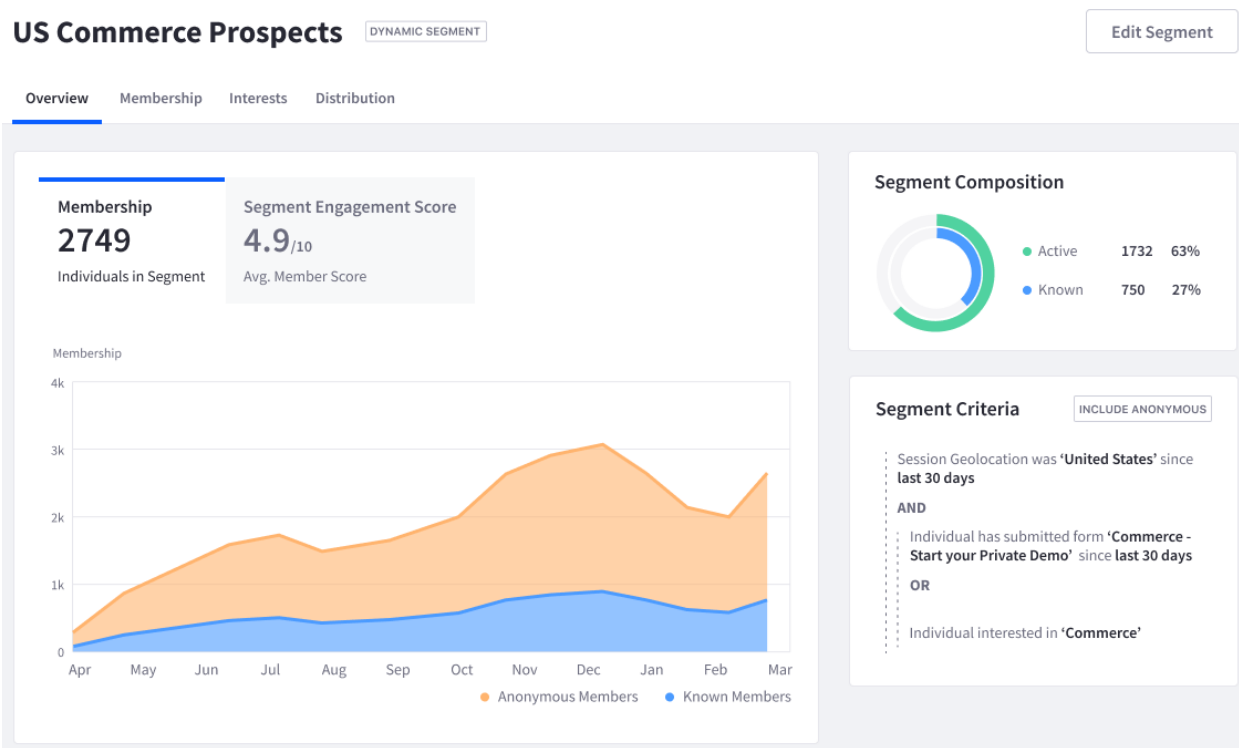

Figure 81.6: The Overview tab's main area shows the Segment's membership growth.

Distribution

The *Distribution* tab breaks down the Segment by a demographic property.

For example, select a property such as *Job Title* to dissect your Segment further. A bar chart appears, displaying a distribution of the Segment member population by Job Title. As a marketer, you could use this information to identify a subset of the Segment to focus a marketing campaign on.

You've completed touring Segment profiles. Creating Segments and analyzing their profiles is key to learning more about your customers.

An in-depth look at the interest score insight is next.

81.5 Customer Insights

Finding out topics people are interested in helps you relate your products and services to them. Analytics Cloud's Interests insight empowers you with rich, easy-to-understand visuals that provide this information.

Interests

As an Individual interacts with the pages and assets, Analytics Cloud notes the topics they contain. The more an Individual visits pages and assets that contain a topic, the higher the topic's ranking is for that Individual's Segments.

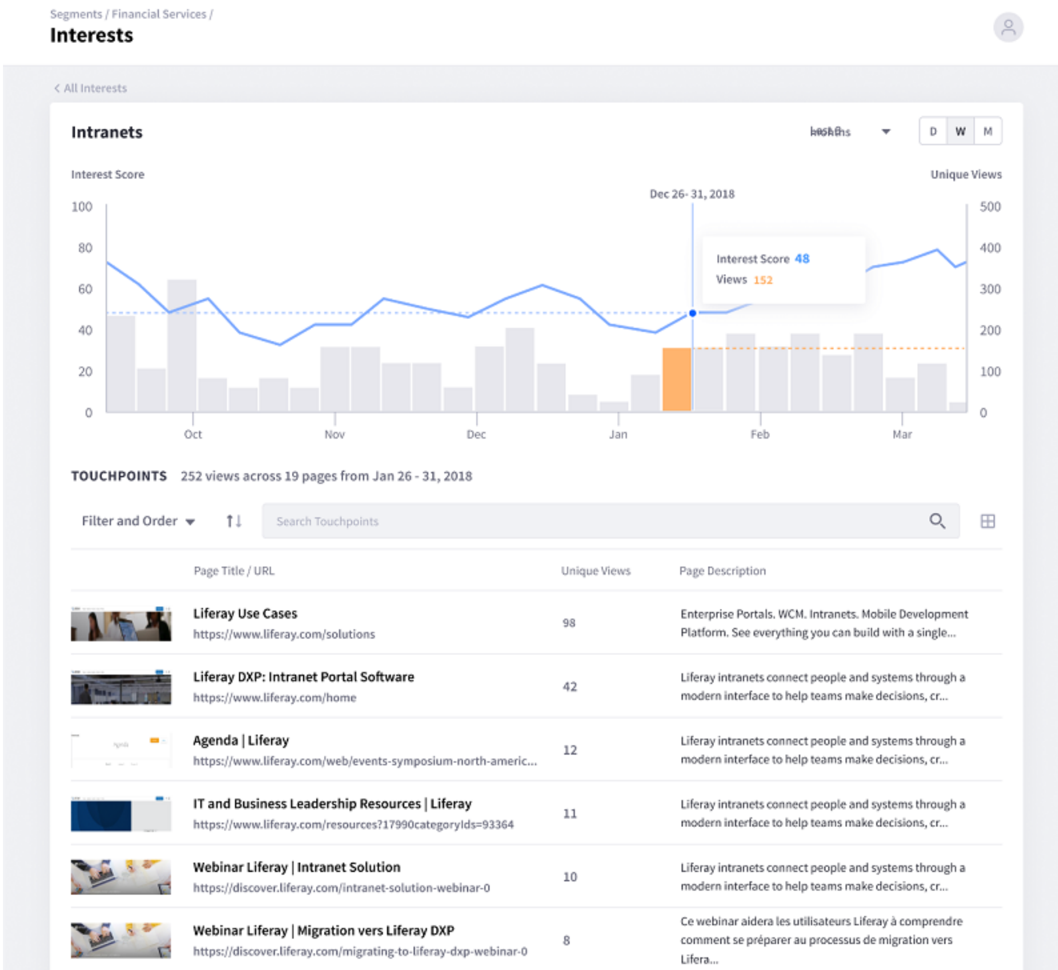

Figure 81.7: Clicking on an Interest topic for a Segment (or Individual) shows a histogram of the Segment's interaction with pages and assets mentioning the topic. In this figure, *Intranets* is the selected topic. Clicking on a histogram point shows the pages and assets viewed during that time period.

For example, say your site has pages that mention these topics:

- Clothing
- Sports

If an Individual named Michael frequents these pages, the Interest ranking for the *clothing* and *sports* topics is high. Here are a couple of different ways you could use this information:

- If you sell sports apparel, you could prepare a campaign for selling sports apparel to Michael.
- If you don't sell sports apparel, you could use the Interest data to convince sports apparel vendors to advertise on your site.

Either way, the Interests insight helps generate new business. The Interests insight tells you the most common topics mentioned by the pages and assets customers visit most.

81.6 Profiling Accounts

If you have a Salesforce data source, Analytics Cloud can import and analyze Salesforce account data. By combining Salesforce account data with data from other sources, Analytics Cloud presents a holistic view of the account and those in it.

To view account analytics, select *Accounts* from the left navigation panel. The accounts are displayed in a searchable table that lists the following data for each account:

- Account Type
- Individuals
- Total Activities
- 30-Day Engagement Score

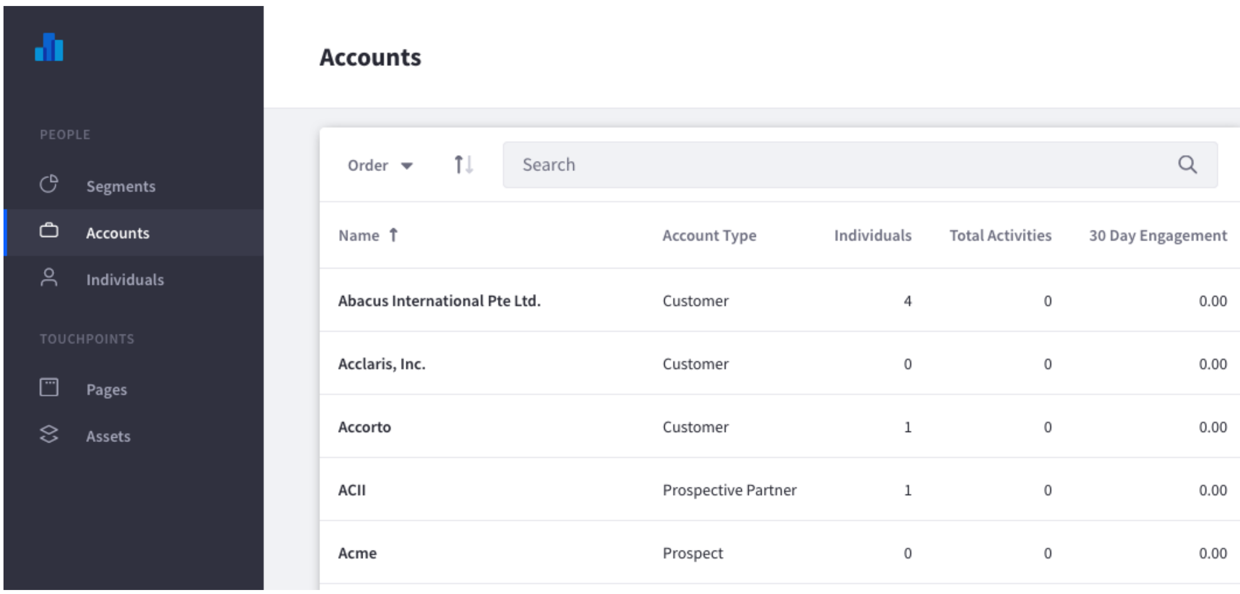

Figure 81.8: The Accounts tab lists the Salesforce accounts imported in Analytics Cloud.

Click an account to view its information, which is organized into these tabs:

- Overview
- Activities
- Interests
- Segments
- Individuals
- Details

The sections that follow describe each tab.

Overview

The Overview tab presents summaries of the account's data. It displays these summaries via these panes:

- Account Activities
- Account Firmographics
- Contact Information
- Account Interest Topics
- Known Individuals
- Associated Segments

If more information is available for each tab, you can access it via its link. For example, a *View All Activities* link is at the bottom of the Account Activities pane. Clicking this link takes you to the Activities tab, which shows more detailed information about activities.

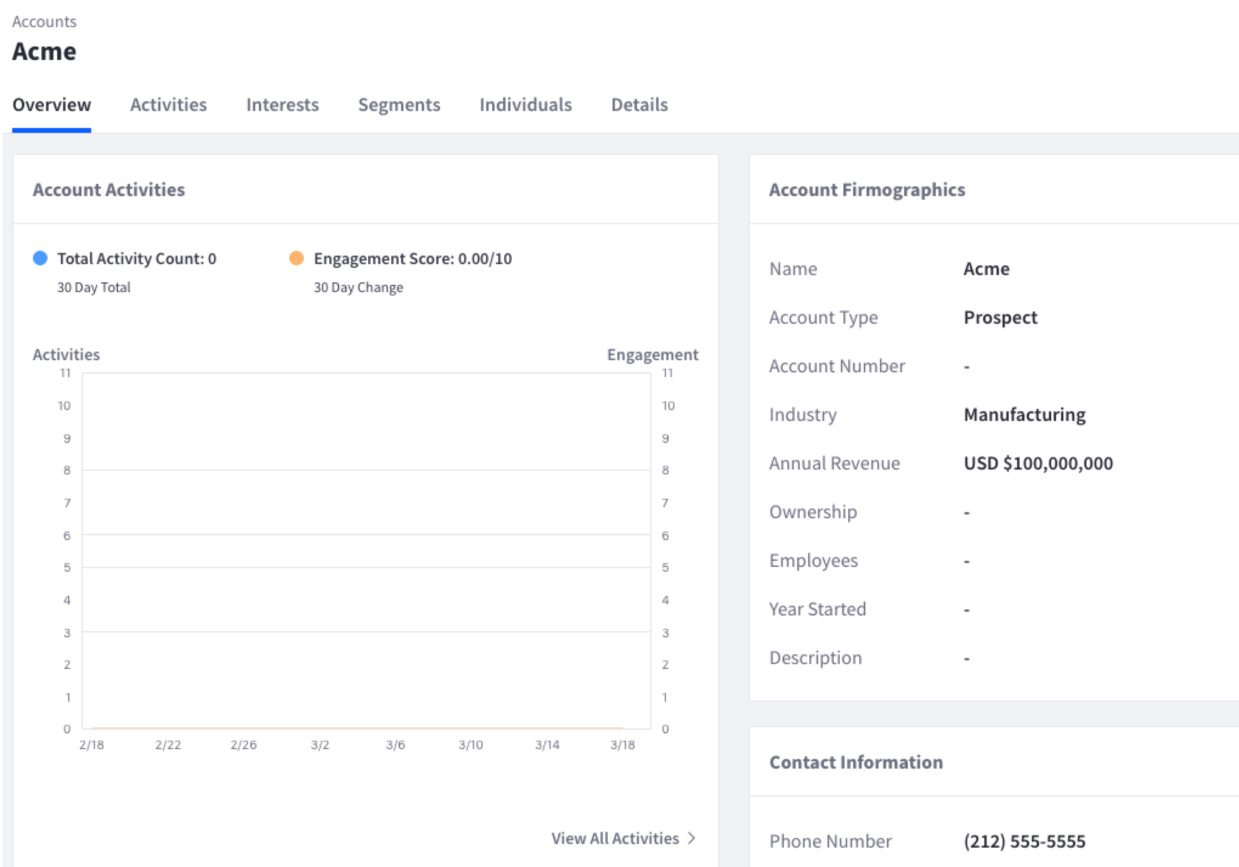

Figure 81.9: The Overview tab displays summaries of the account's data.

Activities

The Activities tab shows the activities of the individuals in the account. A histogram shows the number of activities by date and the average engagement score for the account's individuals. A searchable table below the histogram shows more detail about the activities.

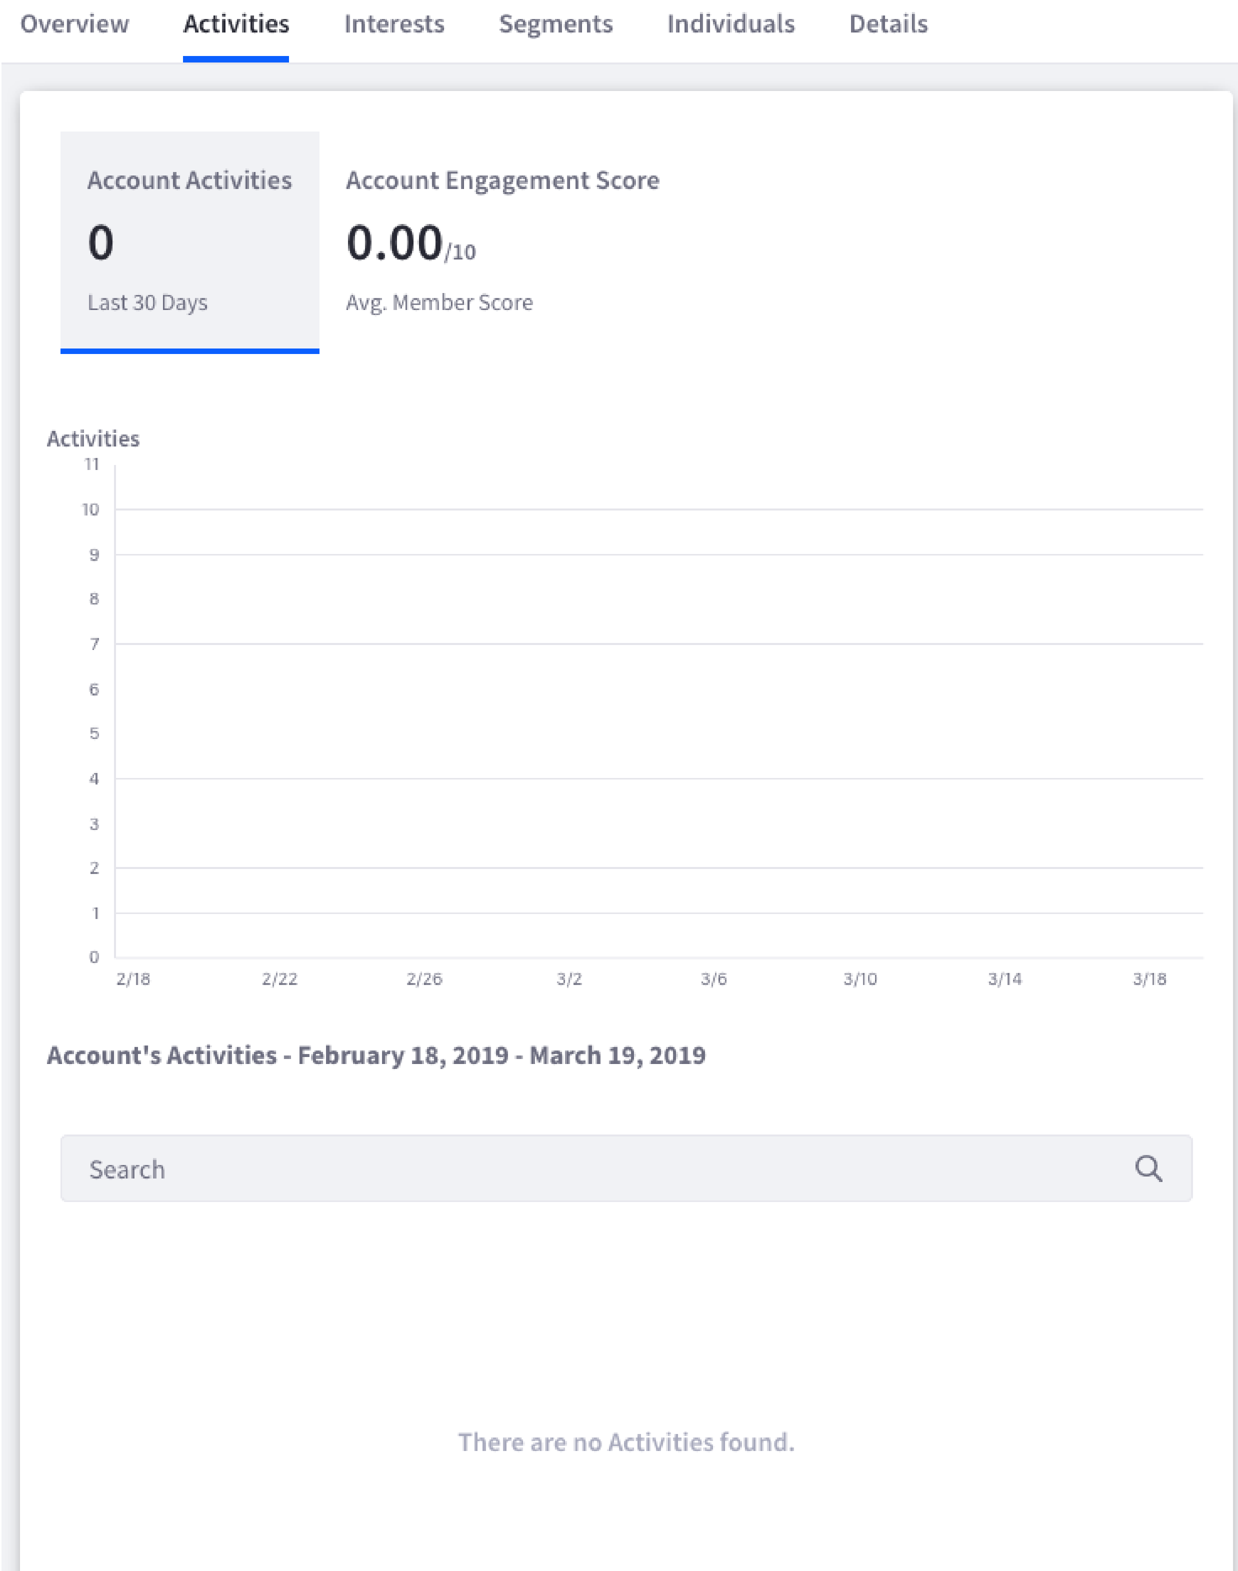

Figure 81.10: The Activities tab displays data on activities of the account's individuals.

Interests

The Interests tab shows the topics of interest for the account's individuals. Analytics Cloud identifies these topics via the same methodology it uses for individuals. For more information, see Customer Insights.

Segments

The Segments tab shows the segments for the account's individuals. These segments function the same as segments in general. The only difference is that the segments here apply to the individuals in the account. For more information on segments, see the documentation on creating and profiling segments.

Individuals

The Individuals tab shows information about the individuals in the account. For more information on individuals in Analytics Cloud, see Profiling Individuals.

Details

The Details tab shows information about the account's properties. The properties appear in a searchable table that shows each account property in Analytics Cloud and its value. The table also has columns for these values:

Source Name: The corresponding property name in the Salesforce data source. For example, the Analytics Cloud `accountId` property is `id` in Salesforce. Therefore, `id` appears in the Source Name column of the `accountId` row.

Data Source: The data source the property's value originated from.

Last Modified: The date the property's value was most recently changed.

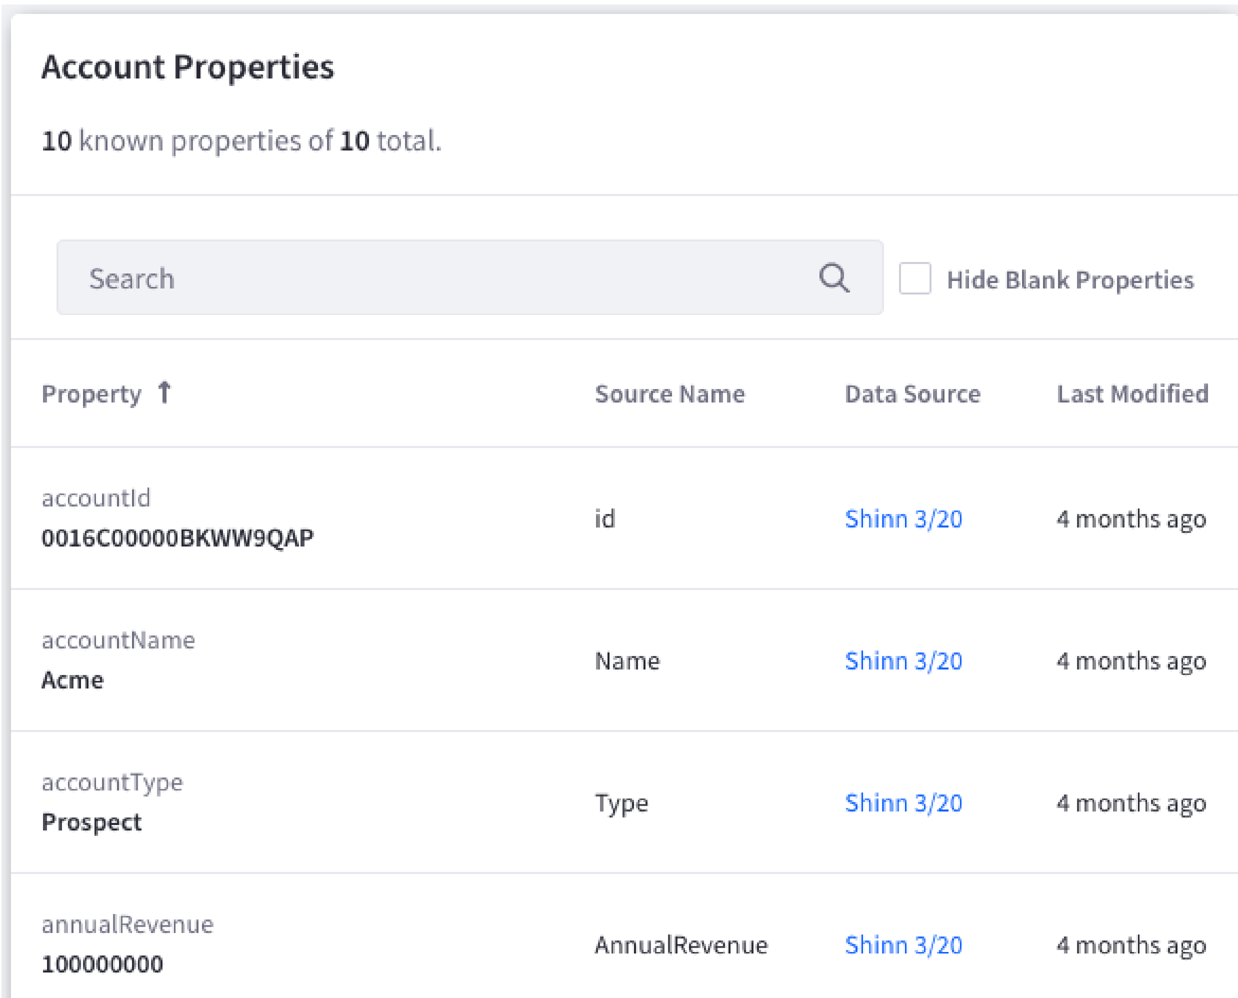

Figure 81.11: The Details tab contains a searchable table that shows the account's properties and their values.

Chapter 82

Finding Analytics Data

Combining traditional page analytics with site, path, and asset analytics sets Liferay Analytics Cloud apart from other analytics tools. Analytics Cloud's site-wide report feature provides a comprehensive report of your entire site. This helps you understand how many unique visitors your site has over time, which pages these visitors access, what their interests are, and more.

Analytics cloud also provides analytics for individual pages. This includes views, visitors, and bounce rate. Path analytics let you see how visitors arrive at your pages. This includes information on which pages they visit prior to yours, and their location and device type. In addition, several Liferay DXP Assets have even more fine-grained metrics to drill deeper than the page level to determine how you can improve your site's performance.

Here, you'll learn how to find the site, page, and asset analytics:

- Site Report
- Viewing Page Data
- Finding Asset Data

The rest of the articles in this section explain these metrics in detail.

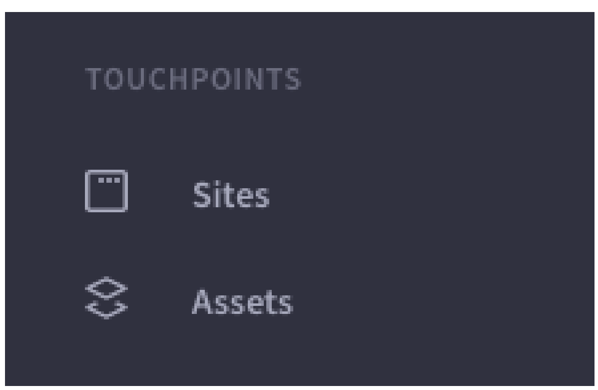

Figure 82.1: Analytics Cloud provides important data for Sites and Assets.

82.1 Site Report

Liferay Analytics Cloud provides a single Site report for each connected data source. For example, if you connect Analytics Cloud to one DXP instance and configure analytics for two subsites, one Site report aggregating both subsites is created.

Follow these steps to view the Site report:

1. In the *Touchpoints* section of the menu, click *Sites*.
2. Click the *Overview* tab (this tab is selected by default when you click *Sites*). This tab contains the Site report.

The Site report contains the following data:

- Site Metrics
- Top Pages
- Acquisitions
- Visitors by Day and Time
- Search Terms and Interests
- Sessions by Location
- Session Technology
- Cohort Analysis

Site Metrics

The Site Metrics panel presents a summary of how visitors interact with your Site. This panel contains the following metrics:

- **Visitors:** Total unique visitors.
- **Sessions per Visitor:** An average of the number of sessions for each unique visitor. A single user can open multiple sessions. A session ends after 30 minutes of inactivity, or at midnight.
- **Session Duration:** The length of time an average session lasts.
- **Bounce Rate:** The percentage of visitors who view your Site's first page, but do nothing else before the session ends.
- **Engagement:** A proprietary metric that Analytics Cloud calculates to determine how engaged a visitor is with your Site. This is derived from session duration, pages visited, scroll depth, time on page, and more.

Clicking each metric changes the visualization in the panel to display the selected metric.

Top Pages

The Top Pages panel lists the following:

- **Visited Pages:** Your Site's most visited pages and the number of visitors for each.
- **Entrance Pages:** The most common pages that visitors view first upon entering your Site, and the number of entrances for each.
- **Exit Pages:** The most common pages that visitors view when leaving your Site, and the exit percentage for each. The exit percentage is the percentage of visitors for which the page is the last page in their session.

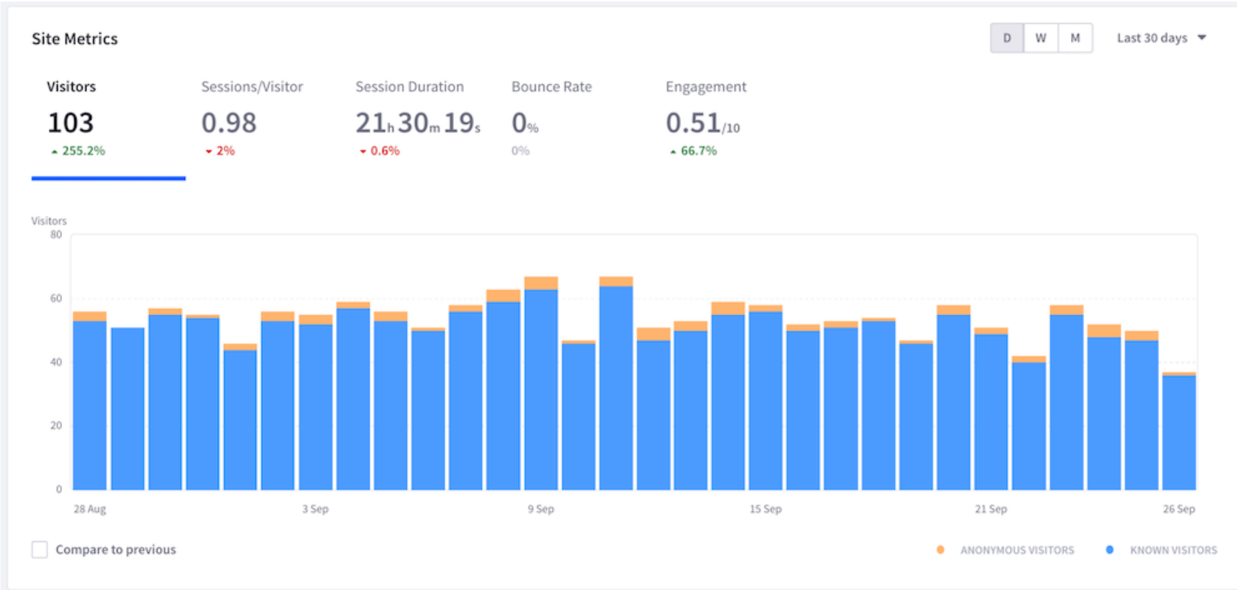

Figure 82.2: The Site Metrics panel shows how visitors are using your site.

To view more detailed page analytics, click *View Pages* in the panel. Alternatively, click the *Pages* tab at the top of the screen. Both take you to the page analytics discussed in Viewing Page Data.

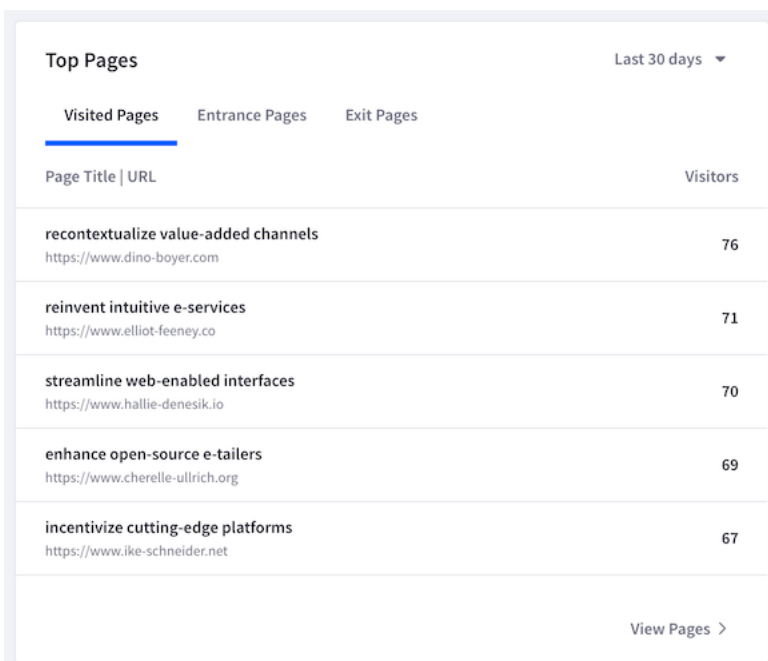

Figure 82.3: The Top Pages panel lists your Site's most relevant pages.

Acquisitions

The Acquisitions panel shows how visitors arrive to your Site. It shows data for the channels, source/medium, and referrers from which visitors arrive to your Site.

This report works best if your marketing campaign utilizes UTM parameters. UTM parameters allow Analytics Cloud to determine where visitors arrive from (e.g., the specific referrer or ad campaign).

Figure 82.4: The Acquisitions panel shows how visitors get to your Site.

Visitors by Day and Time

The Visitors by Day and Time panel visualizes the days and times when visitors come to your Site. This helps you understand when your Site is most active. You can use this information, for example, to know when to release important information or launch an advertising campaign.

The panel contains a grid with the days of the week on one axis, and the time of day on the other axis. Darker cells in the grid indicate heavier Site traffic at the corresponding day and time. Tooltips for each cell show the number of visitors for that day and time.

Search Terms and Interests

The Search Terms and Interests panels show your visitors' most common search terms and the topics they're interested in, respectively. Search terms are collected from the search query parameter in your Site's URL. Interest topics are derived from the keyword metadata of the pages that visitors view. To view the full list of interest topics, click *All interests* in the Interests panel. Alternatively, click the *Interests* tab at the top of the screen.

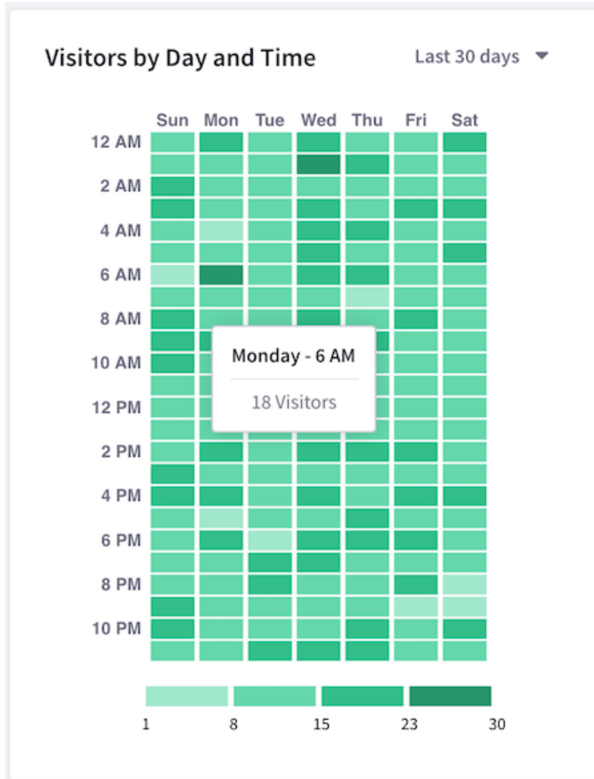

Figure 82.5: The Visitors by Day and Time panel shows your Site's traffic by day and time.

Sessions by Location

The Sessions by Location panel shows the countries from which visitors access your Site. Countries with more visitors are shaded darker on the map. The countries are also listed below the map along with the number and percentage of visitors for each.

Session Technology

The Session Technology panel shows the devices, operating systems, and browsers that visitors access your Site with. Tooltips for each graph element display more detailed data for that element. On the *Devices* tab, for example, mouse over each bar on the bar graph to see the operating system data for that device.

Cohort Analysis

The Cohort Analysis panel shows a cohort analysis based on visitors from a specific acquisition date (the cohort) and whether they return to your Site over a given time period. Use the selector menus at the top left of the panel to select the visitor type (All, Anonymous, or Known) and time period (Day, Week, or Month).

For example, if you select *All Visitors* and *Day*, the percentage of visitors from a given acquisition date that return to your Site are listed in the chart for each following day.

82.2 Viewing Page Data

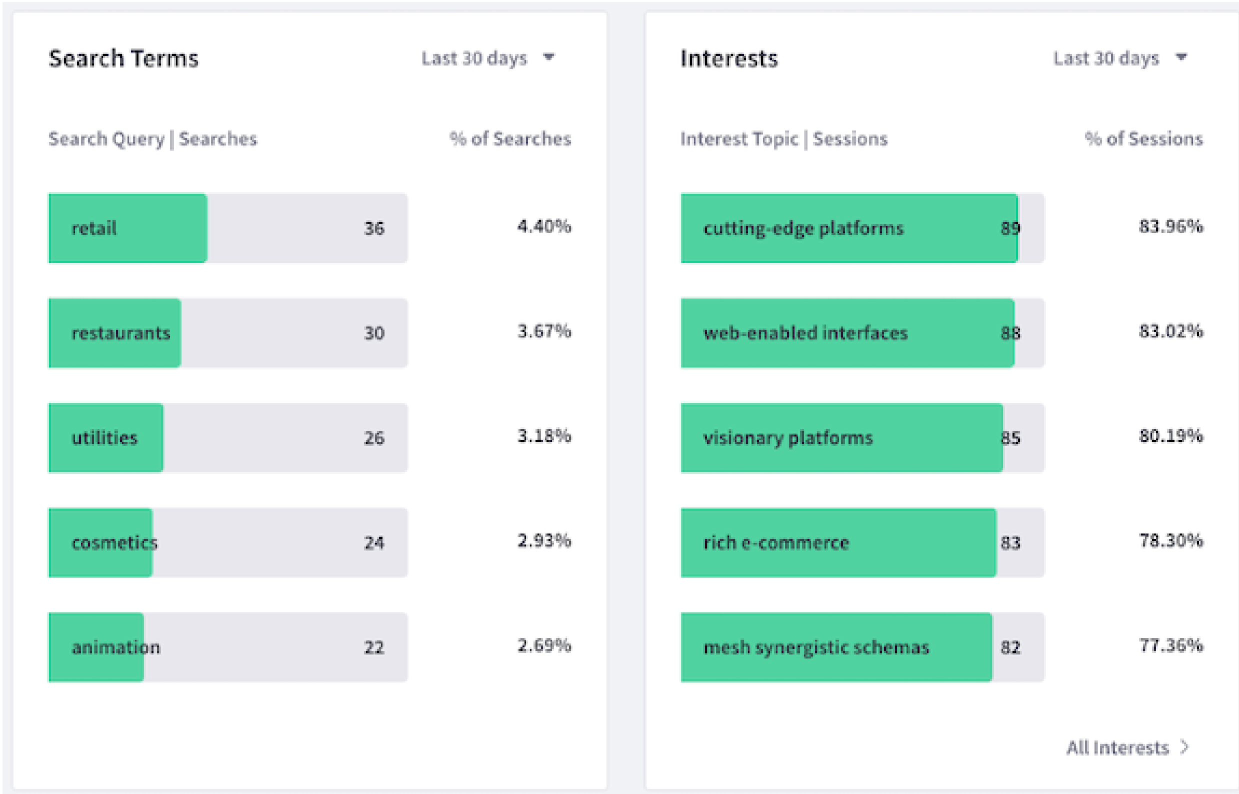

Figure 82.6: The Search Terms and Interests panels show your Site's most common search terms and interest topics.

Follow these steps to view the list of Pages with summary data:

1. Find the Touchpoints section of the menu.
2. Click *Sites*.
3. Click the *Pages* tab.

Order Pages in ascending or descending order based on any one of these metrics. Click the metric's heading in the table to perform the sort:

- Average Engagement Score
- Total Visitors
- Total Views
- Average Bounce Rate
- Average Time on Page

The metrics for ordering the Page list are calculated based on the time period selected in the time period menu (at the top-right of the table). The following values are supported:

- Last 24 hours

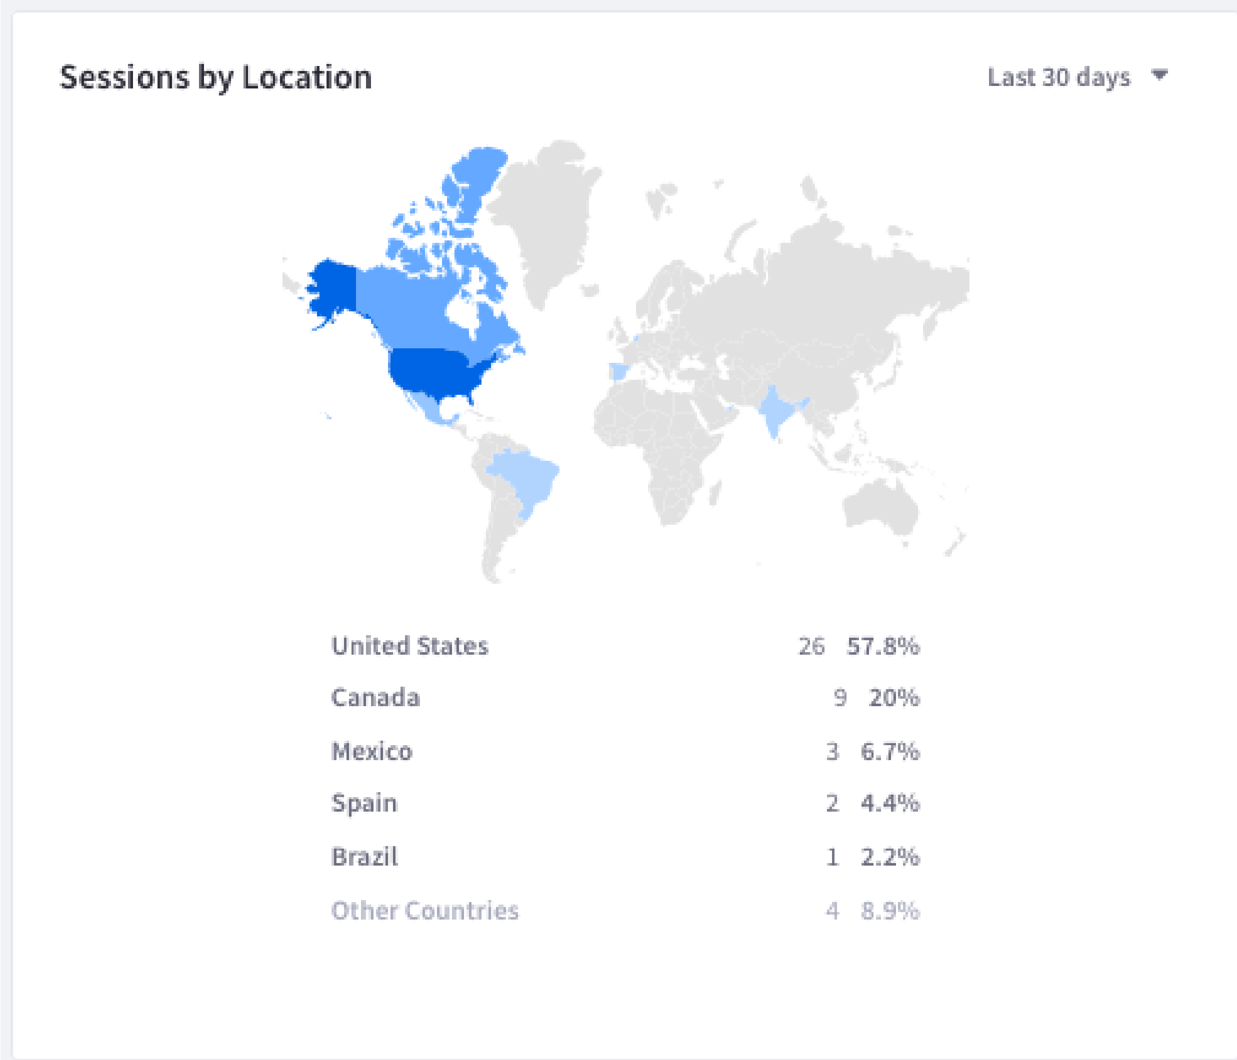

Figure 82.7: The Sessions by Location panel shows where your Site's visitors are.

- Yesterday
- Last 7 days
- Last 28 days
- Last 30 days (default)
- Last 90 days

To view detailed metrics for a single page, click that page in the table. See Page Analytics for more information.

82.3 Finding Asset Data

1. Find the Touchpoints section of the menu.

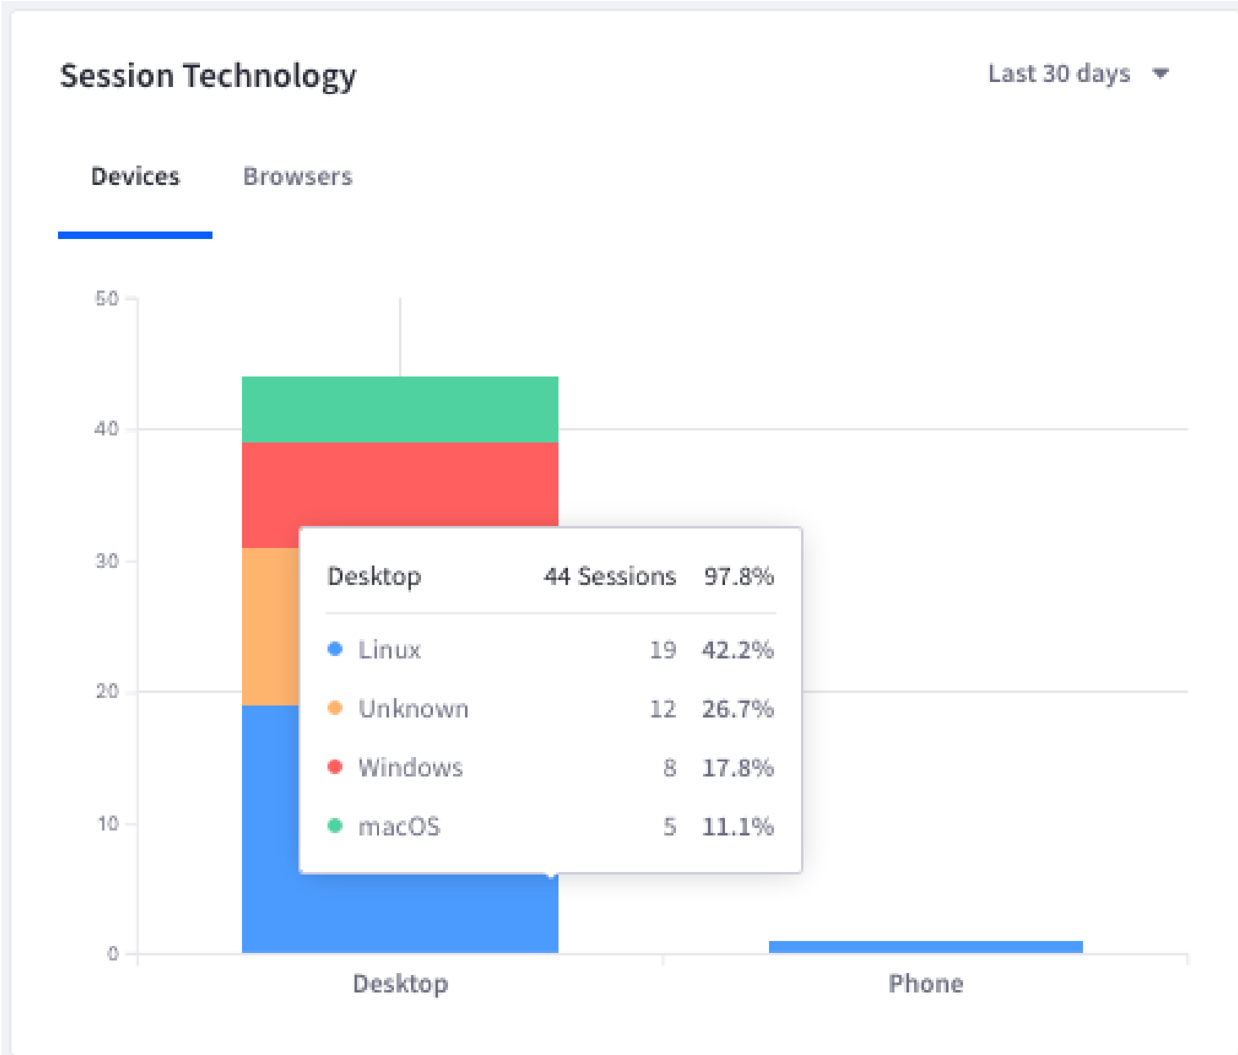

Figure 82.8: The Session Technology panel shows what visitors use to access your Site.

2. Click *Assets* for a list of Assets with summary data.
3. Choose an Asset type: Blogs, Documents and Media, Forms, Web Content, or Custom.

To order Assets in ascending or descending order based on any one of the Asset's metrics, click the metric's heading in the table. The metrics differ between Assets. See each Asset's documentation for more information. In addition, you can calculate the metrics over a selected time period by selecting one of the following values from the time period menu (at the top-right of the table):

- Last 24 hours
- Yesterday
- Last 7 days
- Last 28 days
- Last 30 days (default)
- Last 90 days

Figure 82.9: The cohort analysis is based on visitors from a specific date and the percentage of those visitors that return over the time period that follows.

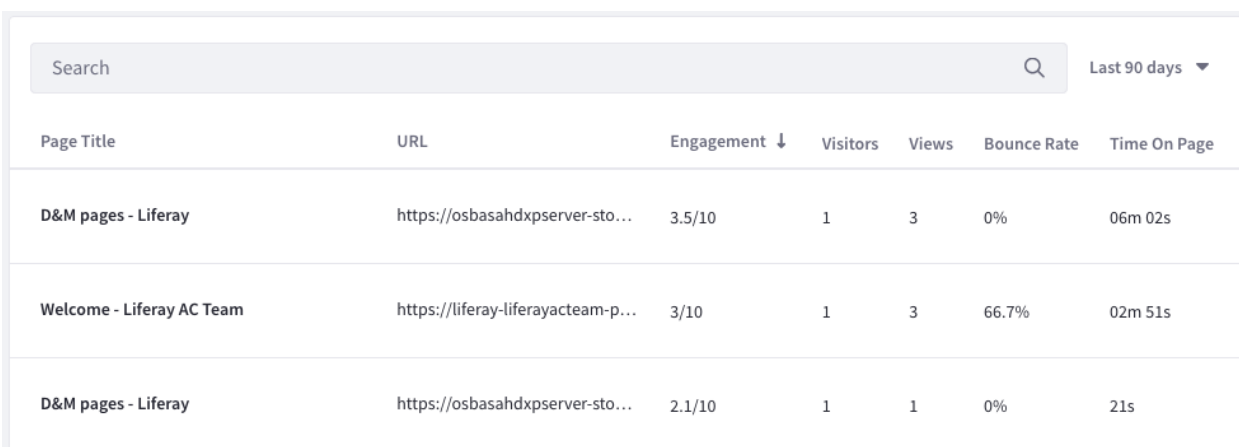

Figure 82.10: : The Page list contains useful summary data.

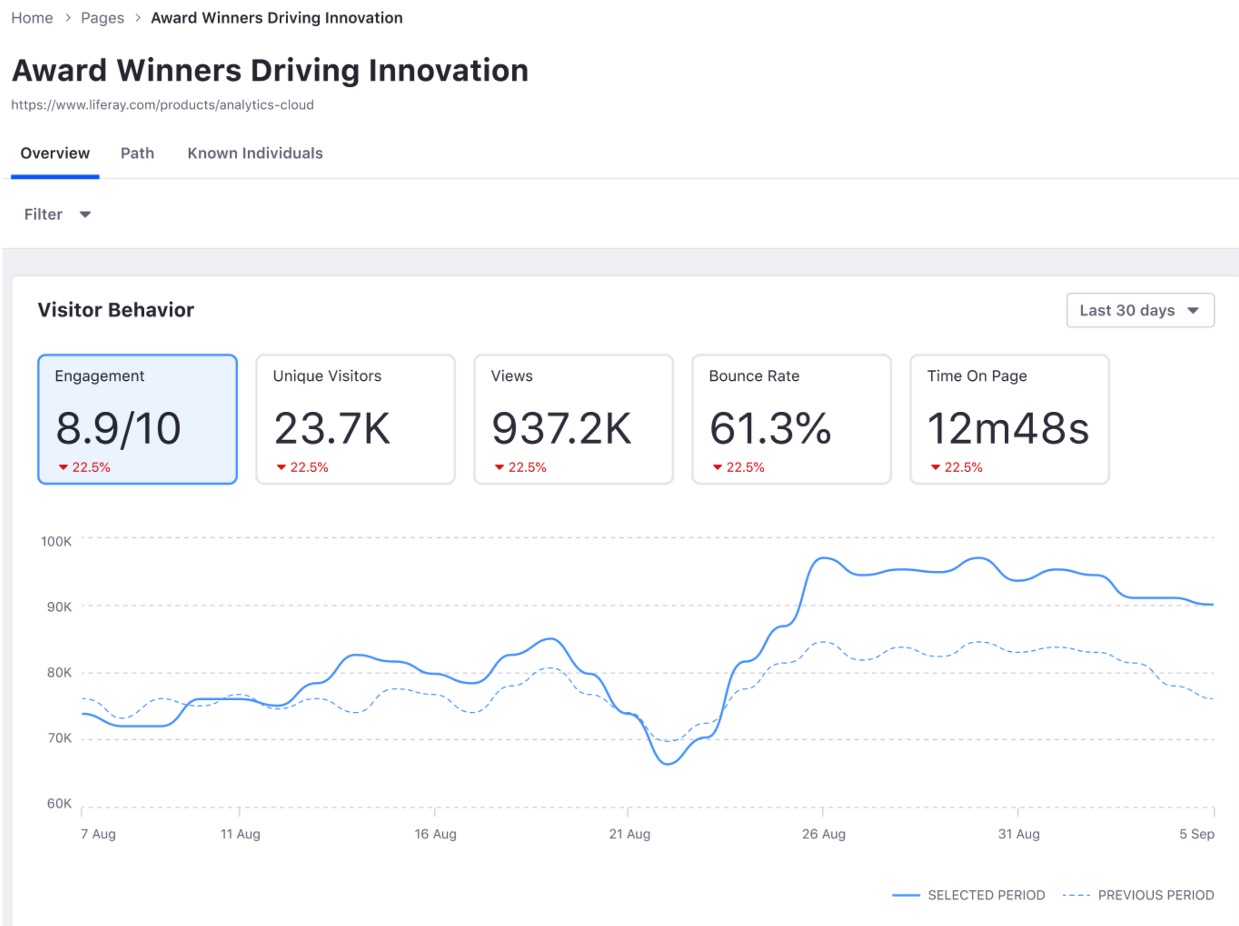

Figure 82.11: : Most Page metrics are available in the Overview tab.

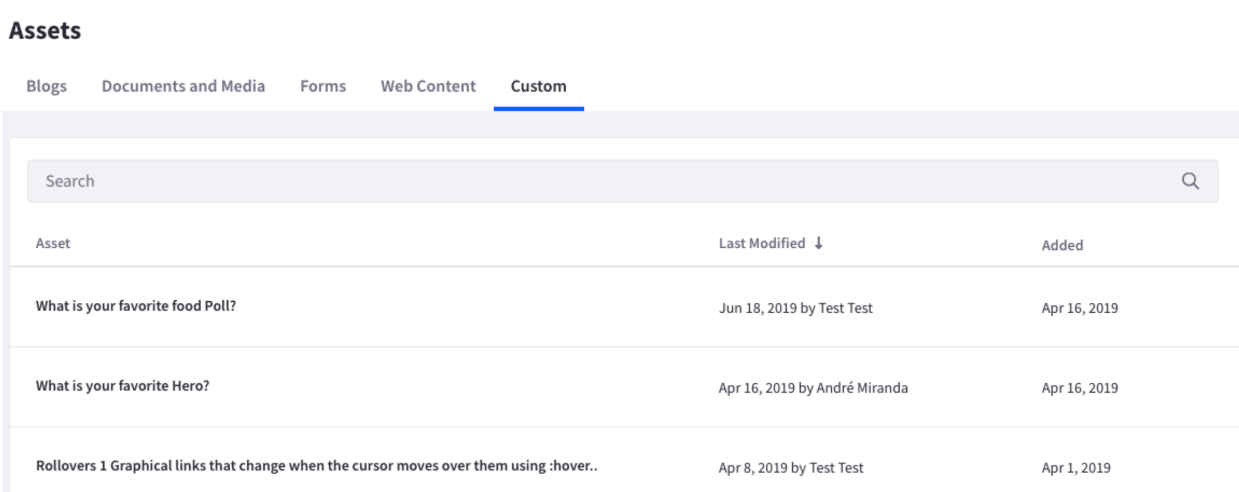

Figure 82.12: : The Assets appear in a table.

Click an Asset in the list to see more granular detail. Remember that there are separate lists for different Asset types. More details on each Asset type's metrics are covered in the subsequent articles.

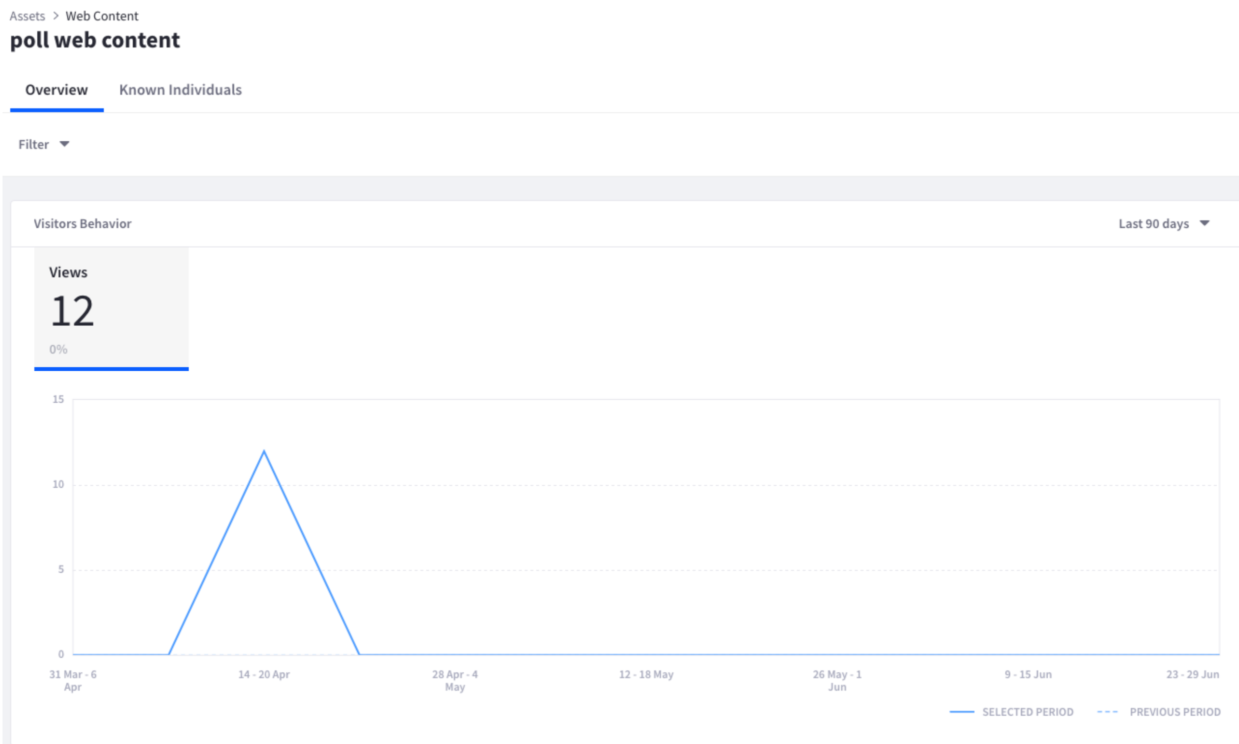

Figure 82.13: : Once you find an Asset, click it to see its metrics.

82.4 Understanding Page Analytics

How are your Site Pages performing? Are visitors abandoning your site when they hit a certain Page? Is one Page constantly getting views, engaging users, and drawing visitors back to your site? Liferay Analytics Cloud answers these questions and more.

Analytics Cloud does the following for pages:

- Registers and tracks pages for analysis the first time the Analytics Cloud client detects a Page interaction.
- Queries pages on a schedule for up-to-date data.
- Reports on pages with the metrics described here.

For instructions on accessing the page analytics discussed here, see Viewing Page Data.

Data Time Periods

All Page data in Analytics Cloud appears for a specified time period. The time period selector recalculates the metrics for the selected time period. You must specify a time period to view the metrics. The default is 30 days, but this is configurable. The following values are supported:

Last 24 hours: Displays data generated over the last 24 hours.

Yesterday: Displays all data generated yesterday, beginning at 12:00 AM and ending at 11:59 PM.

Last 7 days: Displays data generated for the previous seven full days (the current day is excluded), in daily increments.

Last 28 days: Displays data generated for the previous 28 full days (the current day is excluded), in weekly increments.

Last 30 days (default): Displays data generated for the previous 30 full days (the current day is excluded), in weekly increments.

Last 90 days: Displays data generated over a 13 week period beginning 90 days prior to the end of the current week. The time period always begins on a Sunday, and ends with the current, incomplete week (unless viewed on a Saturday approaching midnight).

Note that the first Sunday of the 90 day time period is not included.

Visitor Behavior

The Overview tab of a Page's metrics has several data presentations. The first chart, called Visitor Behavior, contains these metrics:

- Average Engagement Score
- Unique Visitors
- Total Views
- Bounce Rate
- Average Time on Page

Select the time period for the data displayed. There are always two trend lines displayed: Selected Period and Previous Period. This facilitates comparisons between time periods.

So what are these metrics in the Visitor Behavior chart?

Metric: Engagement

Engagement, or average Page Engagement is an aggregation of metrics into one overall score. It incorporates these factors:

- Depth of scroll on the Page
- Number of clicks
- Time spent on the Page
- More!

Use the engagement score as a high level view of the Page's overall performance, as compared with other Pages. It might not tell you specifically what's so effective (or weak) about a Page, but it can tell you if the Page is performing as desired over the selected time period.

Baseball Analogy: Are you a baseball fan? If so you might be familiar with the classic individual metrics used to describe a non-pitcher's proficiency: Batting Average, Slugging Percentage, Fielding Percentage, and a few more. Now there are higher level aggregation metrics used to evaluate baseball players, such as BABIP (Batting Average on Balls In Play) and WAR (Wins Above Replacement). Discussion of those metrics is not the point here. Think of the simpler metrics offered by Analytics Cloud (Views, Visitors, Time on Page) as single snapshots into how useful a Page is, like looking at Batting Average and Fielding Percentage to evaluate baseball players. Think

Figure 82.14: The Visitors Behavior chart contains interesting trend lines.

of Page Engagement as an advanced aggregation metric that captures the overall usefulness of a Page, similar to BABIP or WAR in baseball. Maybe the Analytics team can come up with a WAR-like metric next, to compare how valuable a Page is as compared with the average Page. It could be called TEAR (Page Engagement Above Replacement).

Page Engagement is useful to combine with time period filtering and comparative time period features. Comparing the engagement score from different periods is the best way to determine how your Page's performance changes over time.

Metric: Unique Visitors

Useful with the time period filter, Page Visitors is the number of visitors that accessed a Page in a given period of time.

A unique visitor has a unique IP address in Analytics Cloud. Therefore, if the same person views the Page from a different device, it's logged as two unique visitors.

Metric: Page Views

Useful with the time period filter, Page View is the number of views for a Page Page in a given period of time. It's not the same as the number of visitors, because it doesn't try to count only unique IP addresses. Over the last 30 days, one visitor (IP address) could come back to the Page 100 times. That means there are 100 Page Views, but only one visitor (assuming she used the same machine to access the Page each time. However, a unique view isn't logged for a single user unless at least 30 minutes of inactivity on the Page passes before the user interacts with the Page again.

Metric: Page Bounce Rate

Bounce Rate is the percentage of visitors to the Page that navigated away from the site without any page interaction (including scrolling on the page) after the initial page load. It's calculated as a daily rate (percentage per day), and the daily rate trend line is displayed over the selected time period.

Metric: Time on Page

Time on Page calculates the average time spent on a Page for all the Views each day. It's displayed for the selected time period.

This metric is calculated like this for each 24 hour period:

```
(view-1-time + view-2-time + ...) / total-number-views
```

That concludes the Visitors Behavior chart, but see below for more Page data.

Audience

The Audience report uses charts to present information about the audience interaction with the Page. It answers these questions:

- How many users are interacting with my content?
- Of these users, how many are known or anonymous?
- Of the known users who interacted with my content, how many belong to Analytics Cloud Segments?
- Of the users in Segments, what are the top 15 segments?

Here are the charts in the Audience report:

Unique Visitors: A donut chart that presents the percentage of total unique visitors who are known or anonymous. The center of the chart shows the total number of unique visitors. Tooltips on each chart section show the number of users for that section (e.g, the number of known or anonymous users).

Known Individuals: A donut chart that presents the percentage of known individuals who belong or do not belong to one or more Analytics Cloud Segment. The center of the chart shows the total number of known individuals. Tooltips on each chart section show the number of users for that section (e.g, the number of users belonging to one or more Segments). Click the chart title (Known Individuals) to view a list of all the known individuals who comprise the chart's data.

Known Individuals Segments: A bar chart that shows the percentage of known individuals that comprise each Segment. The chart shows a bar for each of the top five Segments, and then aggregates the remaining Segments into the last bar. A tooltip on the last bar shows the values for each of the remaining Segments.

Views by Location

The map in the Views by Location panel shows the number of views by country in the selected time period.

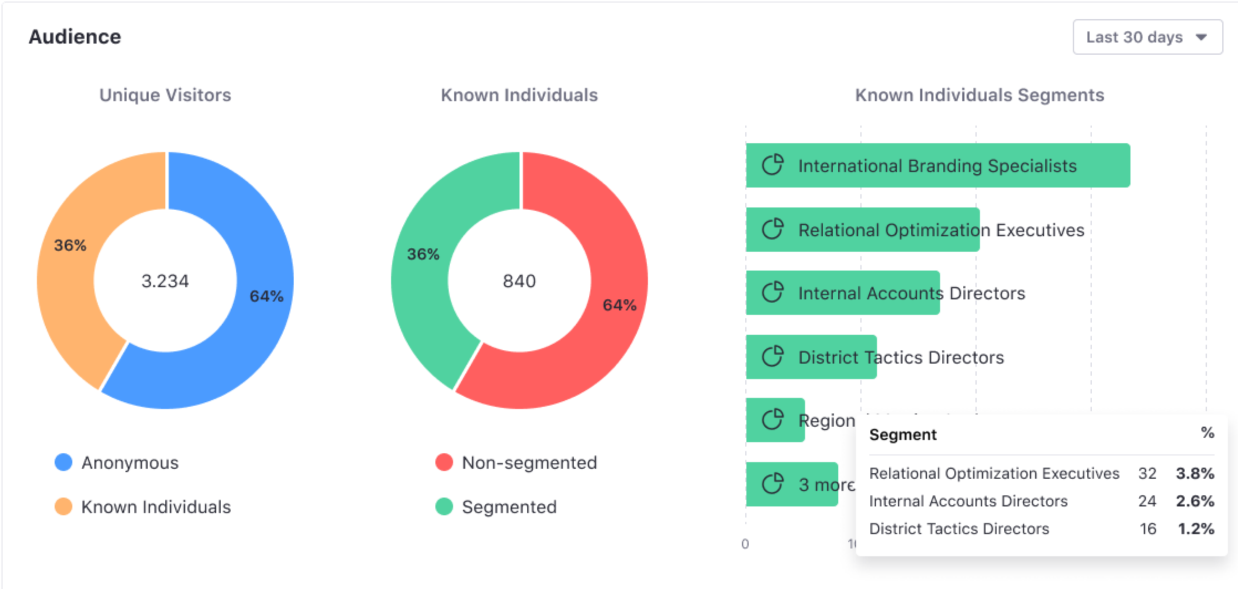

Figure 82.15: The Audience report visualizes how individuals interact with the Page.

Figure 82.16: See where the Page is most popular.

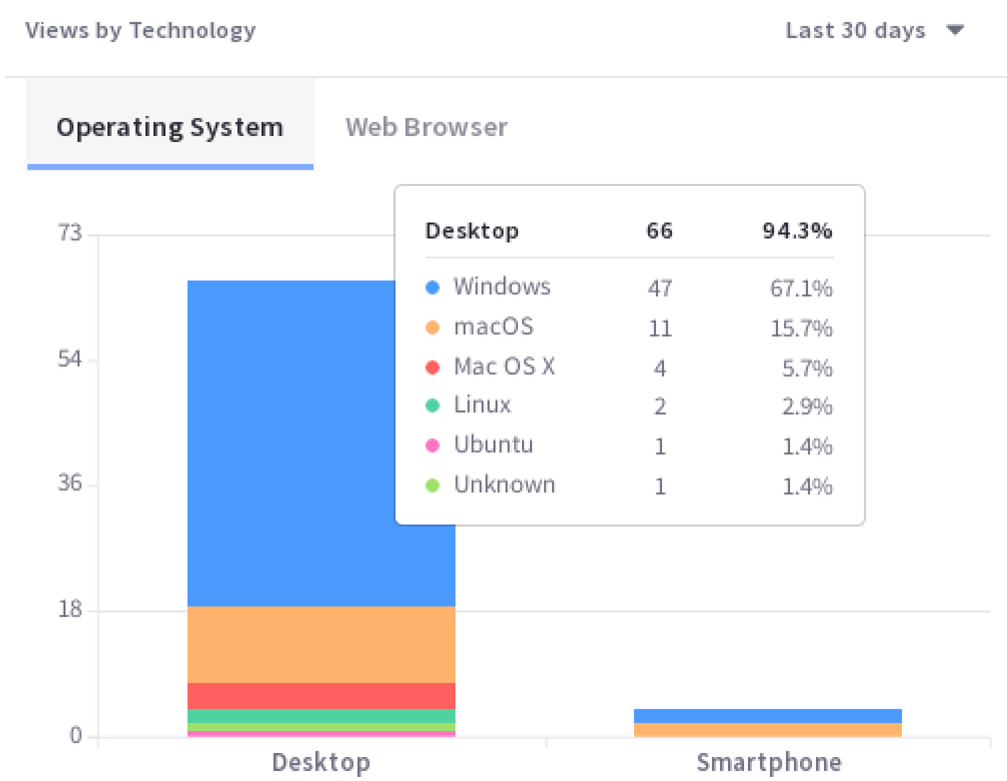

Figure 82.17: Which OS is most commonly being used to access the Page?

Views by Technology

View a stacked bar graph of the Page's views by operating system (grouped by device type) in the default tab. Hover over each bar to see the detailed breakdown of data.

Click *Web Browser* to see a donut chart displaying up to the top eight web browsers over the selected time period. If applicable, remaining web browsers are aggregated in the ninth donut segment.

Assets

View a list of the Assets on the Page by their number of Interactions over the selected time period.

Depending on the Asset being viewed, a different Interaction metric is reported:

- Blogs reports Views.
- Documents and Media reports Downloads.
- Forms reports Submissions.
- Web Content reports Views.

After all those fundamental metrics, you're really getting to know your Pages. But there's some interesting Page data you haven't seen. Discover how people came to the Page in the first place. Learn about Path Analytics next.

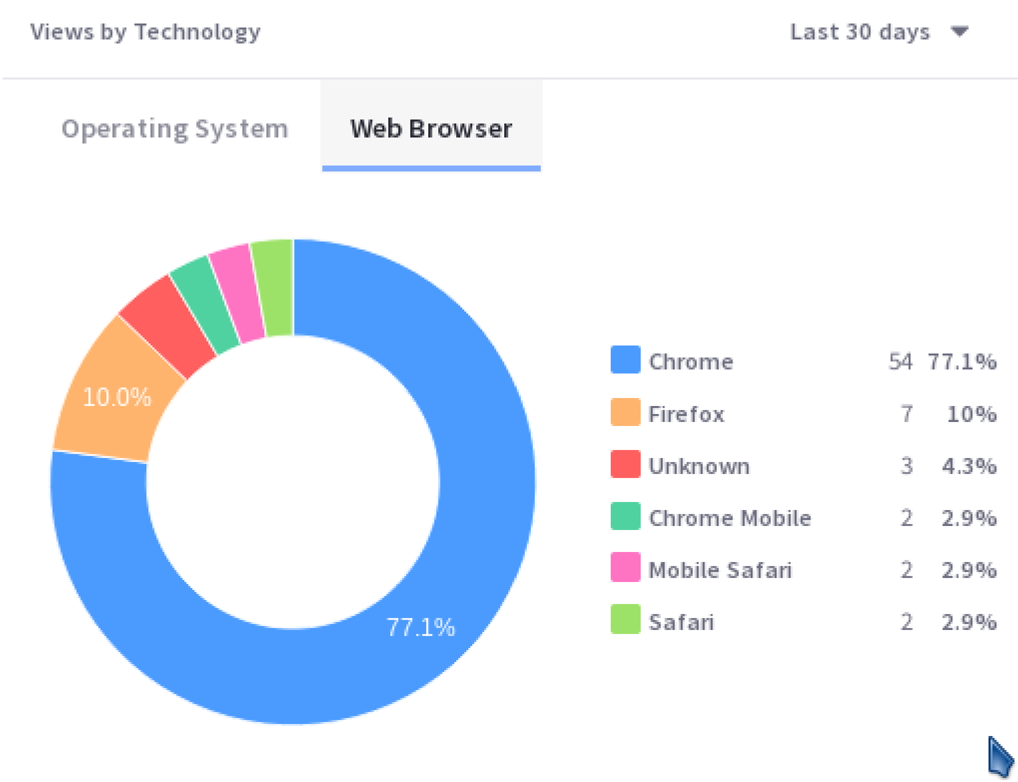

Figure 82.18: Which browser should your page be optimized for?

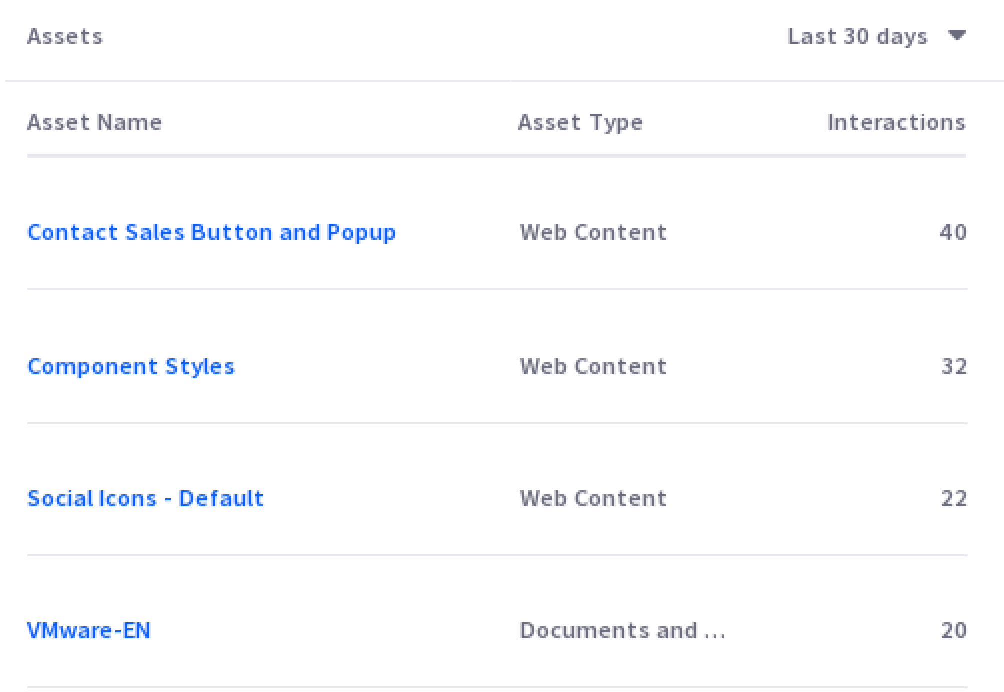

Assets — Last 30 days

Asset Name	Asset Type	Interactions
Contact Sales Button and Popup	Web Content	40
Component Styles	Web Content	32
Social Icons - Default	Web Content	22
VMware-EN	Documents and ...	20

Figure 82.19: Which Assets on the Page are getting the most interactions?

82.5 Path Analytics

A weary sojourner winds his way through the twisted paths of the Internet, seeking answers to life's most important questions. Nearly ready to give up his search and instead look at Grumpy Cat memes, he scrolls to the bottom of the page he's on, skimming the page's content. His weary eyes come to rest on a promising link offered at the very bottom of the page. He clicks it in desperation, hoping it offers the solution to his pressing problem. Through this haphazard placement of an advertisement, he is led to your site, finds exactly the answers he seeks, and becomes a regular visitor or even a customer. If you saw this become a discernible pattern, with a certain page consistently driving visitors to your site, you'd want to target advertising on this page, wouldn't you? Path Analytics identify where the sojourners to your site most commonly originate. Now, instead of hoping they stumble upon your links as you leave them in the gnarled and twisted tubes of the Internet, you can target your placement of links, investing where you're confident they'll be most effective and abandoning fruitless advertising campaigns.

By knowing the Paths your page Visitors use to come to your Page, you'll make better decisions about whether to continue, discontinue, or modify your advertising.

Viewing Page Path Analytics

A Page's Path Analytics are reported in the Path tab of a Page's screen. The diagram reports the most common pages from which visitors accessed this Page.

After the top three paths to the Page, the remaining paths are aggregated to show how many views came from *Other* pages.

If a Page in one of your Liferay DXP sites is also a Path to the current Page, clicking the path URL brings you to its Overview screen. If it's not a page that's loaded from the data source, nothing happens.

Location and Device Type

Path Analytics can be filtered by Location and Device Type. Knowing the Location and Device Type visitors most commonly use to access your site's Pages helps you determine whether your site and its content are optimized properly.

By default, all views of the Page in the selected time period are represented in the Paths diagram. Filter the Paths being viewed by Location and Device Type. Click the *Filter* menu and select one Location whose Path diagram you want to see. The Path diagram is updated to represent the top Paths for only the selected Location. The same filtering can be used for Device Type, and the Device Type and Location filters can be combined. For example, view the Path diagram for views coming from smartphones in Germany.

As you select filters, they're made visible at the top center of the Path screen.

To remove a filter, click the *x* next to the filter name.

Assets

Pages have Assets on them. If a Page containing the Analytics Cloud client is reported in the Path diagram, a *Show Top 5 Assets* link is displayed. Clicking it opens a drop-down list. Up to the top five assets on the Page is displayed, as determined by interaction with the Assets. Each Asset uses the most appropriate interaction metric available:

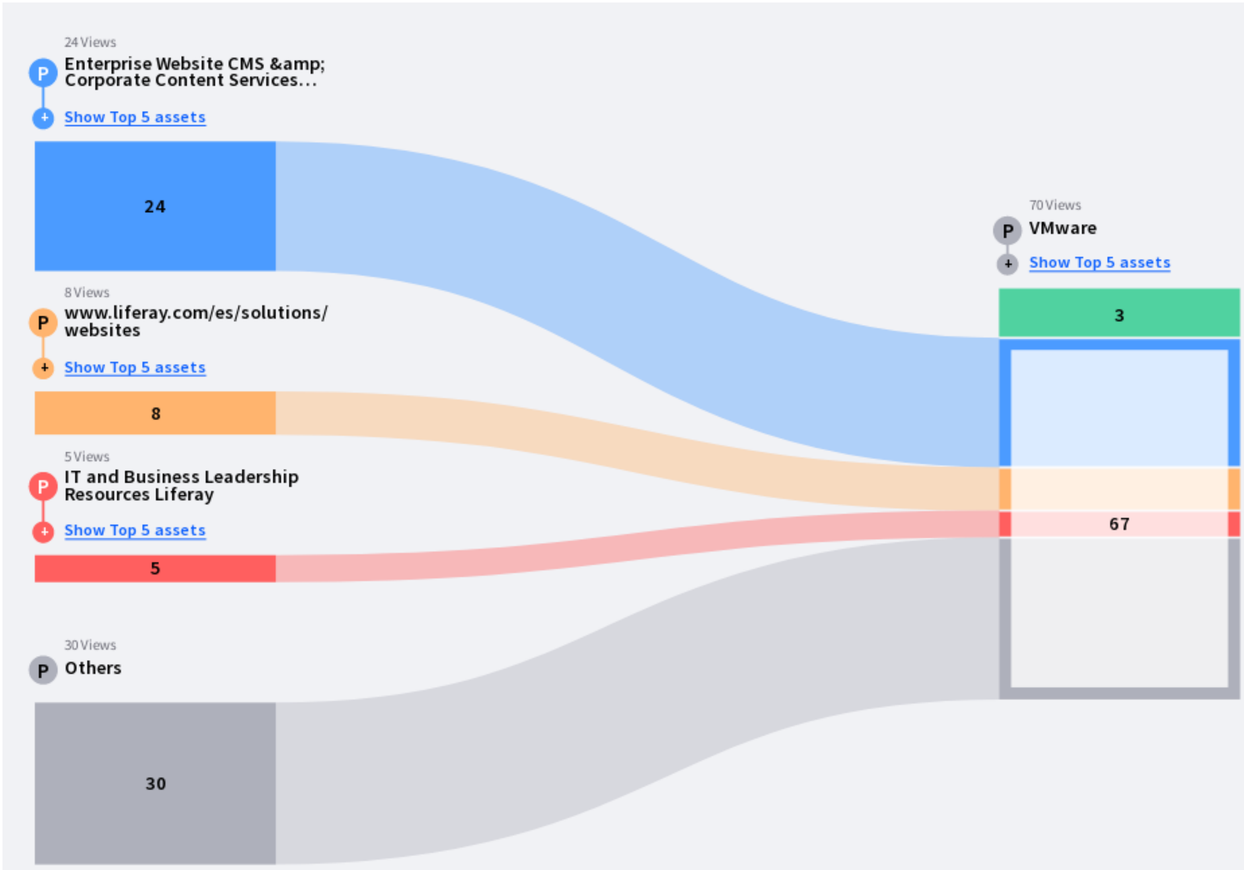

Figure 82.20: On top of being pretty to look at, the Path Diagram contains important information.

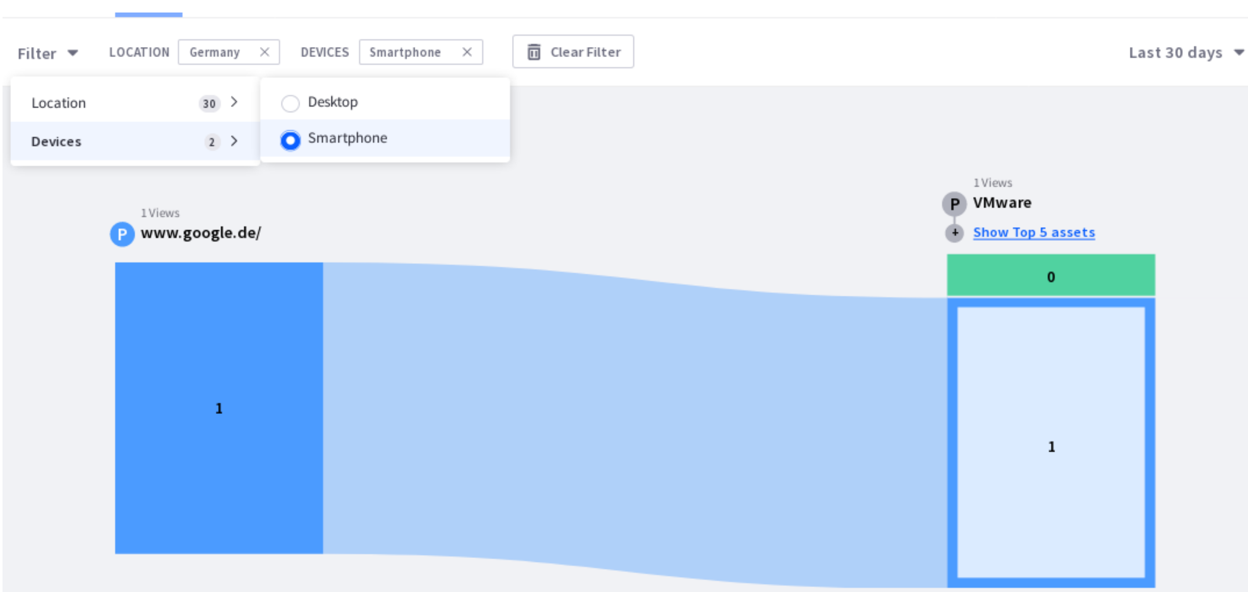

Figure 82.21: One German smartphone User visited this Page in the last 30 days. The User came to the Page from `www.google.com.de`.

- Forms uses Submissions
- Blogs uses Views
- Documents and Media uses Downloads
- Web Content uses Views

Figure 82.22: The top 5 assets are listed for each page.

Clicking an asset brings you to the Asset's Overview page. See the documentation on Assets for more information on analyzing Asset metrics .

82.6 Assets: Analyzing Content

Assets are individual pieces of content that exist on Site Pages in your Liferay Analytics Cloud data sources. In summary, Assets are

- Content residing on a Page.
- Registered, tracked, and reported on by Analytics Cloud the first time an interaction is detected by the Analytics Cloud client.
- Queried and reported on at regular intervals.
- Reported on in Analytics Cloud with the data reports shown in this tutorial.

Metrics for these Assets are currently reported in Analytics Cloud:

- Forms
- Blogs
- Documents and Media

- Web Content

Continue reading for details on how to interpret Analytics Cloud data to better understand your assets' performance.

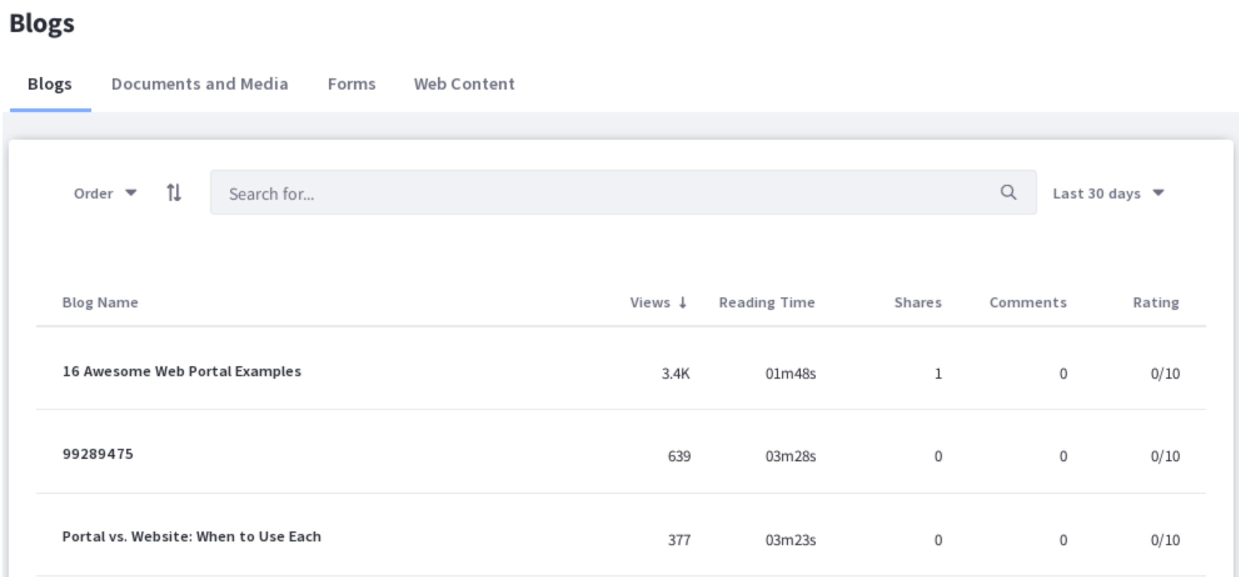

Blog Name	Views	Reading Time	Shares	Comments	Rating
16 Awesome Web Portal Examples	3.4K	01m48s	1	0	0/10
99289475	639	03m28s	0	0	0/10
Portal vs. Website: When to Use Each	377	03m23s	0	0	0/10

Figure 82.23: Each Asset has its own table.

Each Asset is covered in a separate tutorial.

82.7 Form Analytics

Forms are important direct data gathering tools for enterprises. Are your web forms providing you invaluable information or turning users away? Analytics Cloud gives you important insights.

In Analytics Cloud, find the Touchpoints section of the menu and click *Assets*. There are four tabs, each displaying a paginated list of the Asset Type indicated by the tab title. Click *Forms*.

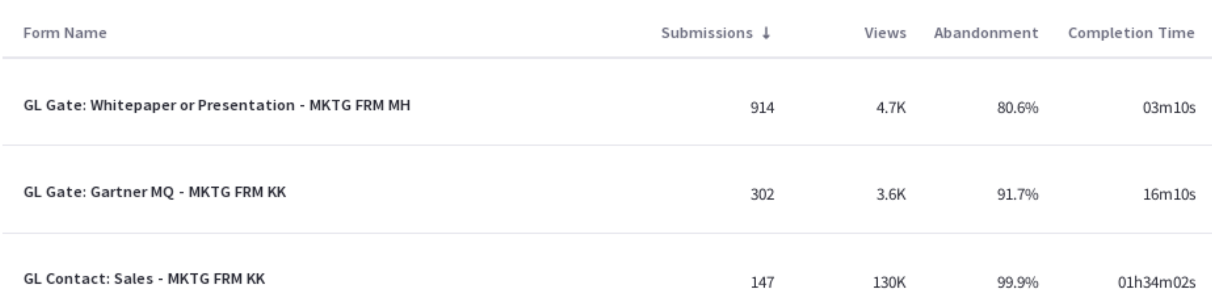

Form Name	Submissions	Views	Abandonment	Completion Time
GL Gate: Whitepaper or Presentation - MKTG FRM MH	914	4.7K	80.6%	03m10s
GL Gate: Gartner MQ - MKTG FRM KK	302	3.6K	91.7%	16m10s
GL Contact: Sales - MKTG FRM KK	147	130K	99.9%	01h34m02s

Figure 82.24: By default, Forms are listed in descending order of Submissions.

View the summary metrics for an Asset directly in the list, and click on an Asset to view its detail page.

Visitor Behavior

As with all Asset types, the Visitors Behavior chart is at the top of its page and provides a line graph with several trend lines. Choose from four important metrics:

- Total Submissions
- Total Views
- Abandonment rate
- Average Completion Time

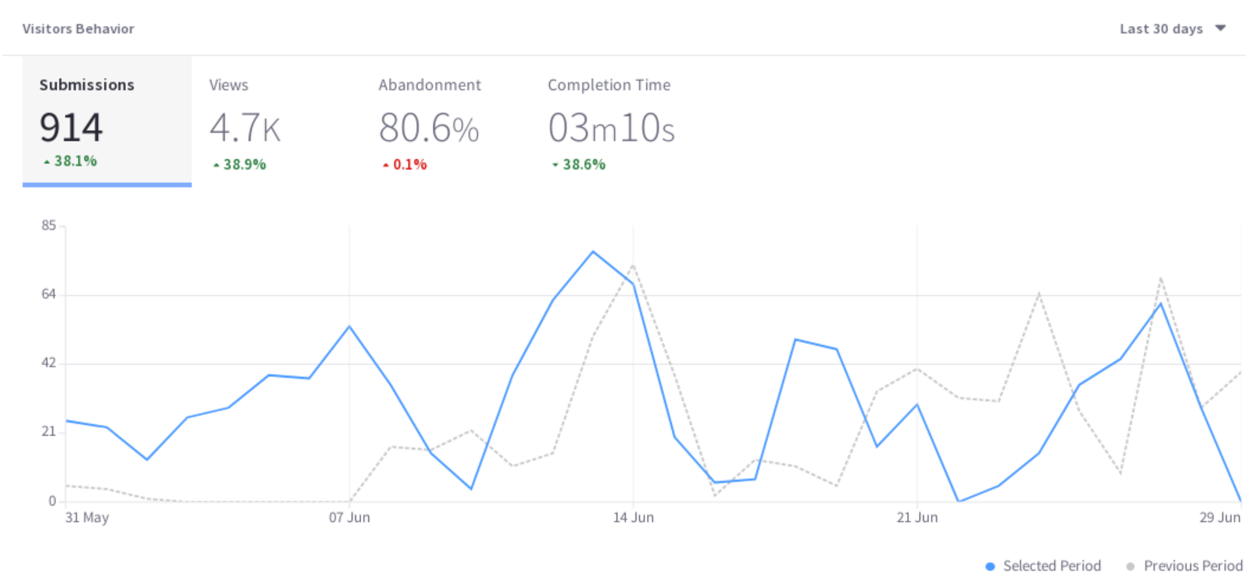

Figure 82.25: The Visitors Behavior chart contains important trend lines.

Submissions

Submissions counts the number of times the Submit button was clicked on a Form. It's the gold standard metric for a form, because that's why the form was created in the first place: to collect data entered into the form. If the Submit button isn't clicked, you don't get the data you wanted.

The Submissions trend line shows the number of times the Submit button was clicked each day (or hour, if *Yesterday* or *Last 24 Hours* are selected) over the selected time period.

Views

Views is a common metric among all Assets (and Pages).

Useful with the time period filter, Views is the number of views for a Form in a given period of time. It's not the same as the number of visitors, because it doesn't try to count only unique IP addresses. So over the last 30 days, one visitor (IP address) could come back to the Page 100 times.

That means there are 100 Page Views, but only one visitor (assuming the same device was used to access the Page each time). However, a unique view isn't logged for a single user unless at least 30 minutes of inactivity on the Page passes before the user interacts with the Page again.

Abandonment

Abandonment is the daily (or hourly if *Yesterday* or *Last 24 Hours* is selected) percentage of users that interacted with the form but stopped short of submitting an entry over the selected time period.

Completion Time

Completion time is a daily average (or hourly if *Yesterday* or *Last 24 Hours* is selected) of the time it took for form users to go from their first interaction with the form until they hit the Submit button.

The time series metrics displayed in the Visitors Behavior chart are paramount to understanding Asset performance over time. But there's more to Asset Analytics.

Audience

The Audience report uses charts to present information about the audience interaction with the asset. It answers these questions:

- How many users are interacting with my content?
- Of these users, how many are known or anonymous?
- Of the known users who interacted with my content, how many belong to Analytics Cloud Segments?
- Of the users in Segments, what are the top 15 segments?

Here are the charts in the Audience report:

Unique Visitors: A donut chart that presents the percentage of total unique visitors who are known or anonymous. The center of the chart shows the total number of unique visitors. Tooltips on each chart section show the number of users for that section (e.g, the number of known or anonymous users).

Known Individuals: A donut chart that presents the percentage of known individuals who belong or do not belong to one or more Analytics Cloud Segment. The center of the chart shows the total number of known individuals. Tooltips on each chart section show the number of users for that section (e.g, the number of users belonging to one or more Segments). Click the chart title (Known Individuals) to view a list of all the known individuals who comprise the chart's data.

Known Individuals Segments: A bar chart that shows the percentage of known individuals that comprise each Segment. The chart shows a bar for each of the top five Segments, and then aggregates the remaining Segments into the last bar. A tooltip on the last bar shows the values for each of the remaining Segments.

Submissions by Location

The map in this panel shows the number of submissions by country in the selected time period.

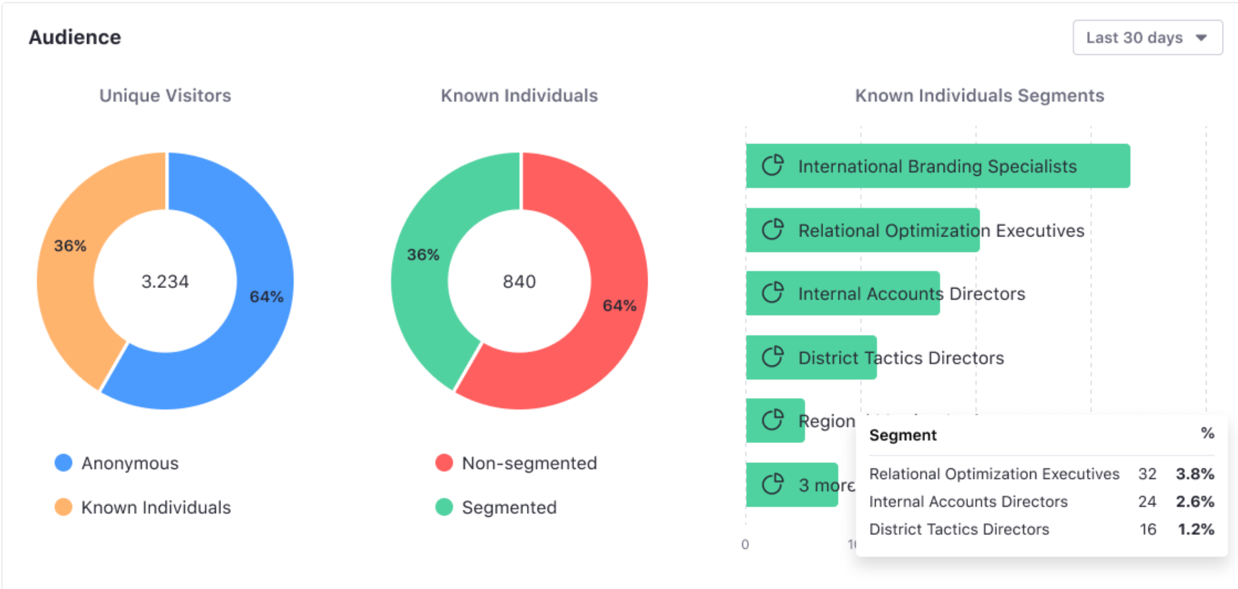

Figure 82.26: The Audience report visualizes how individuals interact with the asset.

Figure 82.27: From which location do users submit the form most frequently?

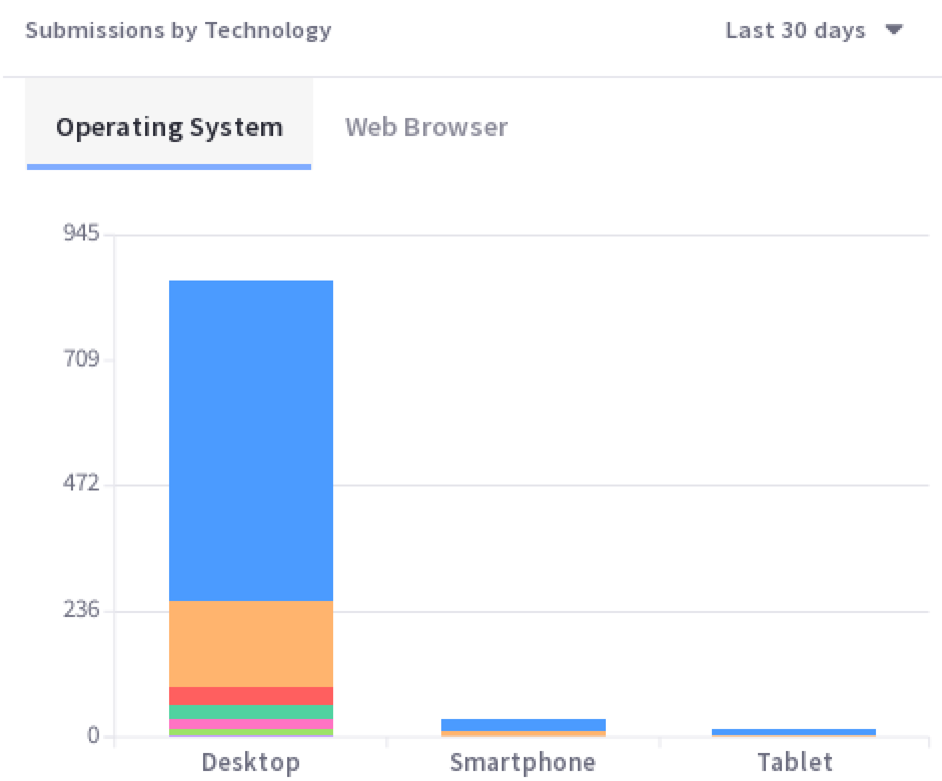

Figure 82.28: What technologies are used to interact with the Asset?

Submissions by Technology

View a stacked bar graph of the Page's submissions by operating system (grouped by device type) in the default tab.

Click *Web Browser* to see a donut chart displaying up to the top eight web browsers over the selected time period. If applicable, remaining web browsers are aggregated in the ninth donut segment.

Field Analysis

The Field Analysis chart is unique to Forms.

Field analysis appears as a bar graph. The height of the first bar shows the total number of people (Views) who have interacted with the form in any way, even just scrolling past it on the page. The rest of the bars show the percentage of those people that interacted with a specific field by clicking into it. The gross numbers on the right y-axis are the number of people who interacted with the field by clicking into it.

Field Analysis gives you these metrics on each form field:

- Number of clicks into the field
- Percentage of Users who abandoned the form at this field.
- The time spent interacting with the field (includes abandonments)

- The refill percent for the field. Out of the total number that filled out the field, this answers the question, "What percentage had to edit or enter new information into the field after moving on in the form?"

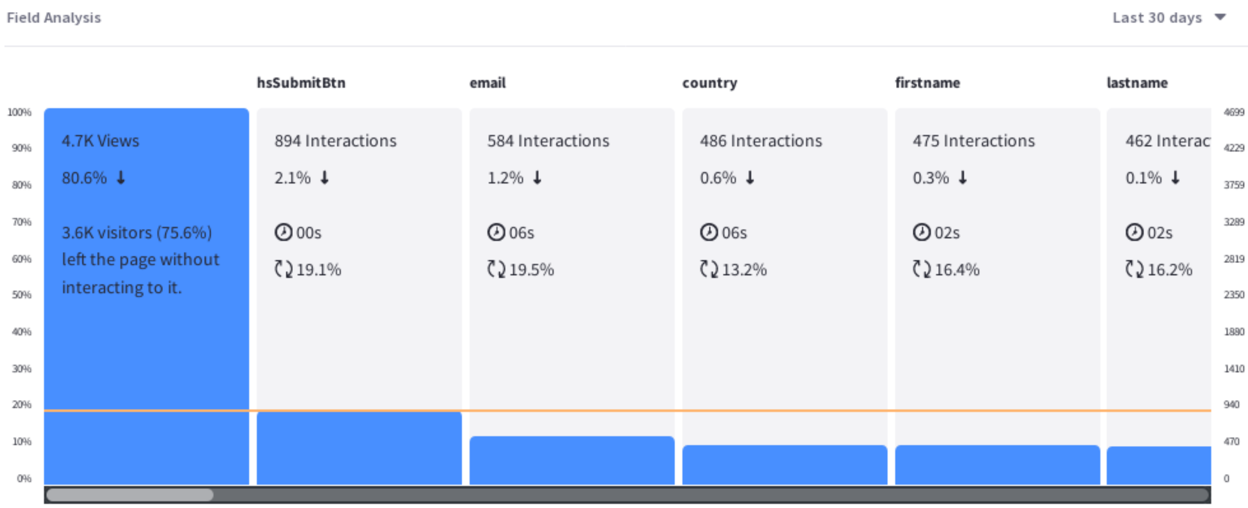

Figure 82.29: See how Form Users interact with the Form's fields.

Field Analysis helps you get down to the field level in diagnosing potential issues with your forms. For example, in your Newsletter Sign-Up form, if you included a required field that asked for users to provide their first child's full name, you may see a lot of form abandonment on that field, not only because some form users don't have children. If you make that field not required, you'll get less abandonment on that field, and if you remove that field from the form entirely, you'll eliminate all abandonment from offended parents.

Asset Appears On

The Asset Appears On table includes a Page Name and URL. Like any Page data in Analytics Cloud, only Pages that were interacted with in some way by site visitors are tracked and reported. It's important to note that the link doesn't take you to the actual Page with the Asset on it. Instead, it leads to the Page Analytics view of the Page. From there you can click the URL to the actual Page.

82.8 Blogs Analytics

Blogs Entries are important content on websites. They attract readers who return to view new blog posts and interact in discussion.

Analytics Cloud gives you important insights into how engaged readers are with your blog posts.

Visitor Behavior

The Overview tab of an Asset's detail presents several data presentations. The first chart, called Visitor Behavior, contains five important metrics:

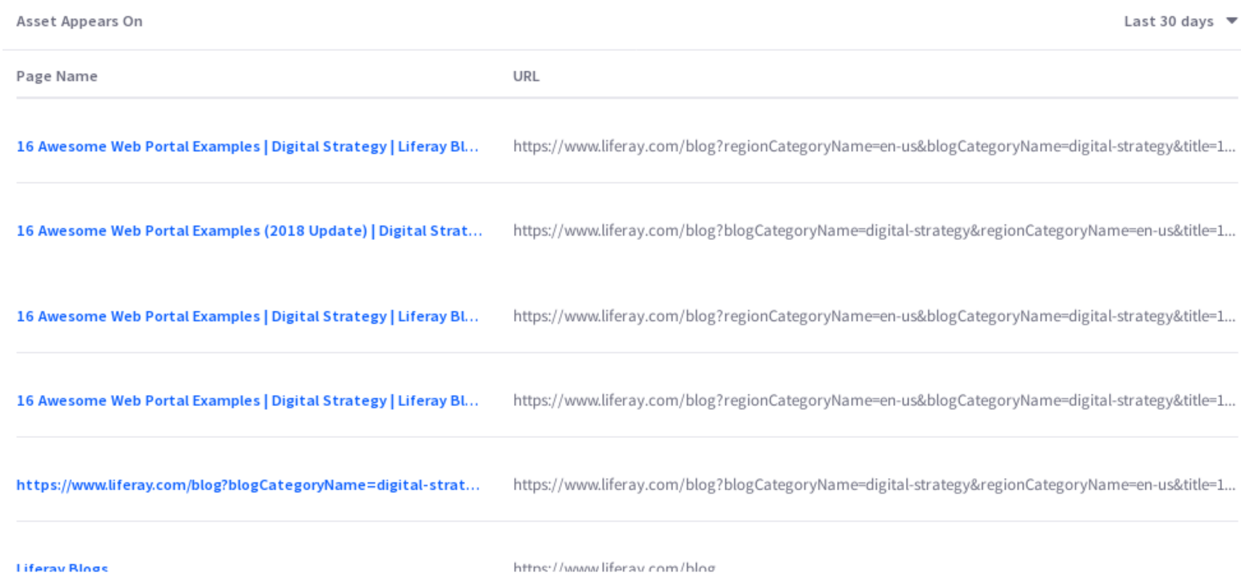

Figure 82.30: What Pages does the Asset appear on?

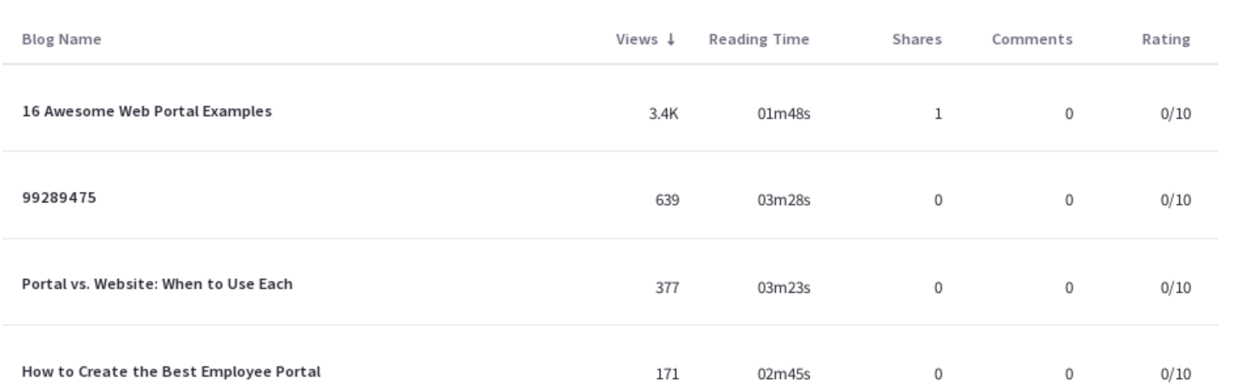

Figure 82.31: By default, Blogs are listed in descending order of Views.

- Total Views
- Average Reading time
- Total Shares
- Total Comments
- Average Rating

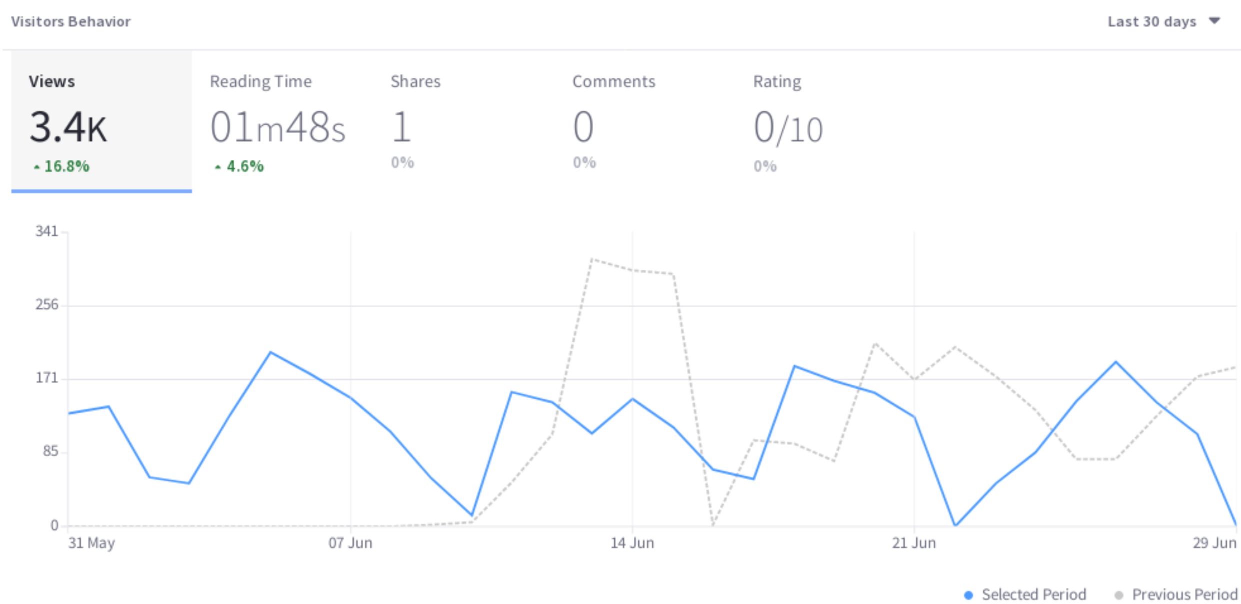

Figure 82.32: The Visitors Behavior chart contains important trend lines.

Views

The Views plot shows the trend-line for the number of times a Blogs Entry was viewed each day (or hour, if *Yesterday* or *Last 24 Hours* are selected).

Reading Time

Reading Time is the average reading time per view of a Blogs Entry. The daily (or hourly, if *Yesterday* or *Last 24 Hours* are selected) average reading time per view is plotted for the time period.

Shares

The count of total Shares per day is plotted to establish the Shares trend-line.

Comments

The Comments trend-line plots the total number of comments on a Blogs Entry per day (or hour, if *Yesterday* or *Last 24 Hours* are selected) over the selected time period.

Rating

Readers can rate blogs.. Whether the rating is a simple Thumb Up/Thumbs Down, or a Stars rating (0-5 Stars), the Rating in Analytics Cloud lets you compare Blog posts to each other. A Blog post announcing Free Swag may have a higher rating than a Blog post announcing that prices are increasing. The rating is reported as a number out of 10.

Audience

The Audience report uses charts to present information about the audience interaction with the asset. It answers these questions:

- How many users are interacting with my content?
- Of these users, how many are known or anonymous?
- Of the known users who interacted with my content, how many belong to Analytics Cloud Segments?
- Of the users in Segments, what are the top 15 segments?

Here are the charts in the Audience report:

Unique Visitors: A donut chart that presents the percentage of total unique visitors who are known or anonymous. The center of the chart shows the total number of unique visitors. Tooltips on each chart section show the number of users for that section (e.g, the number of known or anonymous users).

Known Individuals: A donut chart that presents the percentage of known individuals who belong or do not belong to one or more Analytics Cloud Segment. The center of the chart shows the total number of known individuals. Tooltips on each chart section show the number of users for that section (e.g, the number of users belonging to one or more Segments). Click the chart title (Known Individuals) to view a list of all the known individuals who comprise the chart's data.

Known Individuals Segments: A bar chart that shows the percentage of known individuals that comprise each Segment. The chart shows a bar for each of the top five Segments, and then aggregates the remaining Segments into the last bar. A tooltip on the last bar shows the values for each of the remaining Segments.

Views by Location

The map in the Views by Location panel shows the number of views by country in the selected time period.

Views by Technology

View a stacked bar graph of the Page's views by operating system (grouped by device type) in the default tab.

Click *Web Browser* to see a donut chart displaying up to the top eight web browsers over the selected time period. If applicable, remaining web browsers are aggregated in the ninth donut segment.

Asset Appears On

The Asset Appears On table includes a Page Name and URL. Like any Page data in Analytics Cloud, only Pages that were interacted with in some way by site visitors are tracked and reported. It's

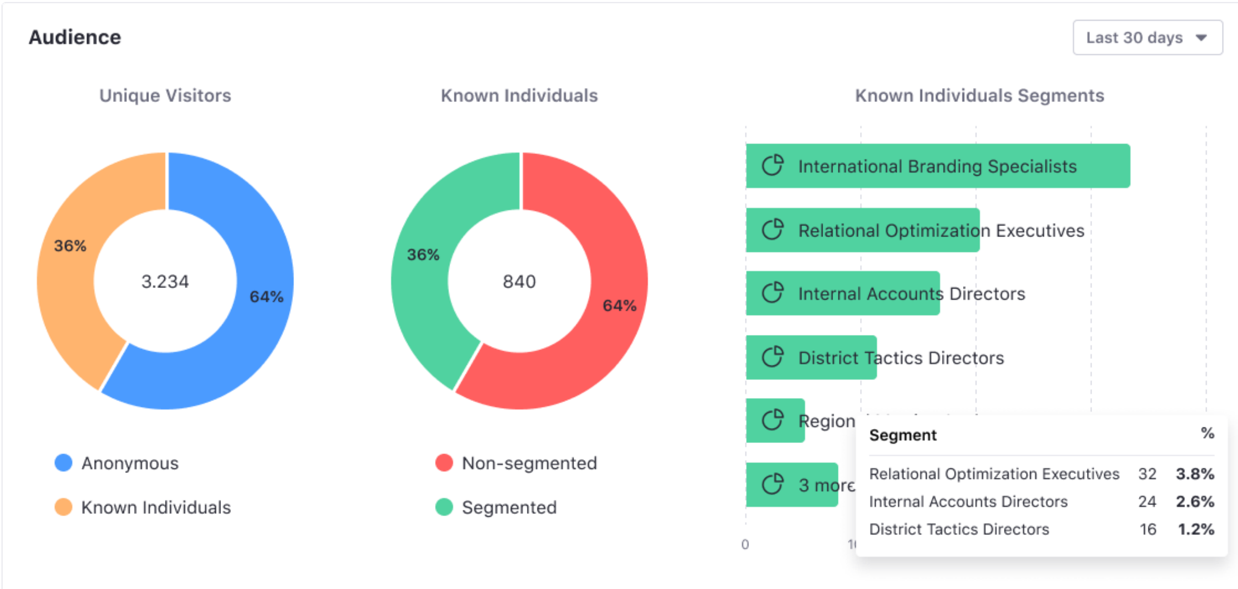

Figure 82.33: The Audience report visualizes how individuals interact with the asset.

Figure 82.34: From which location do users interact with the Asset most frequently?

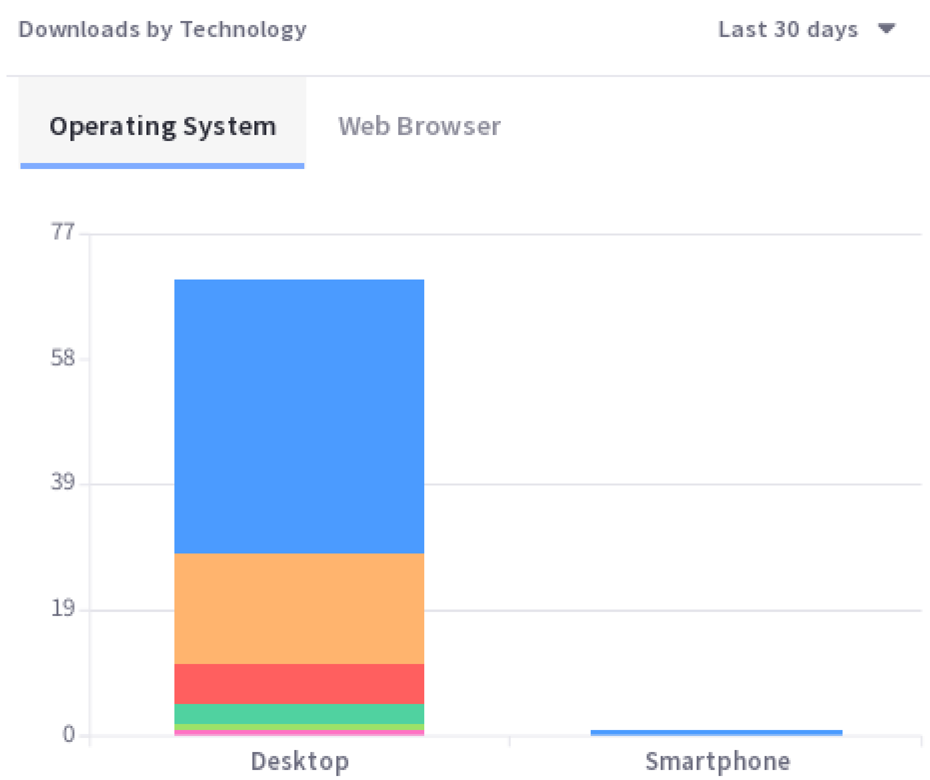

Figure 82.35: What technologies are used to interact with the Asset?

important to note that the link doesn't take you to the actual Page with the Asset on it. Instead, it leads to the Page Analytics view of the Page. From there you can click the URL to the actual Page.

82.9 Documents and Media Analytics

Documents and Media files are assets you want site users to look at or download.

Are your Documents and Media engaging site visitors? Analytics Cloud can give you important insights.

Visitor Behavior

The Overview tab of an Asset's detail presents several data presentations. The first chart, called Visitor Behavior, contains four important metrics:

- Total Downloads
- Total Previews
- Total Comments
- Average Rating

Asset Appears On — Last 30 days

Page Name	URL
16 Awesome Web Portal Examples | Digital Strategy | Liferay Bl...	https://www.liferay.com/blog?regionCategoryName=en-us&blogCategoryName=digital-strategy&title=1...
16 Awesome Web Portal Examples (2018 Update) | Digital Strat...	https://www.liferay.com/blog?blogCategoryName=digital-strategy&regionCategoryName=en-us&title=1...
16 Awesome Web Portal Examples | Digital Strategy | Liferay Bl...	https://www.liferay.com/blog?regionCategoryName=en-us&blogCategoryName=digital-strategy&title=1...
16 Awesome Web Portal Examples | Digital Strategy | Liferay Bl...	https://www.liferay.com/blog?regionCategoryName=en-us&blogCategoryName=digital-strategy&title=1...
https://www.liferay.com/blog?blogCategoryName=digital-strat...	https://www.liferay.com/blog?blogCategoryName=digital-strategy&regionCategoryName=en-us&title=1...
Liferay Blogs	https://www.liferay.com/blog

Figure 82.36: What Pages does the Asset appear on?

Document Name	Downloads ↓	Previews	Comments	Rating
Allianz Case Study	71	0	0	0/10
Coachweb Intranet Case Study	61	0	0	0/10
Volkswagen Group France - EN	52	0	0	0/10

Figure 82.37: By default, Documents and Media files are listed in descending order of Downloads.

Downloads

The Downloads trend-line shows the total number of downloads per day (or hour, if *Yesterday* or *Last 24 Hours* are selected) over the selected time period.

Previews

Documents and Media Documents can be previewed. Even if not ultimately downloaded, previewing shows interaction with the file and is therefore an important metric. The Previews trend line shows the total number of previews per day (or hour, if *Yesterday* or *Last 24 Hours* are selected) for the selected time period.

Figure 82.38: The Visitors Behavior chart contains important trend lines.

Comments

Comments can be allowed on Documents and Media files. If enabled, the Comments trend line shows the total number of Comments on a file per day (or hour, if *Yesterday* or *Last 24 Hours* are selected) over the selected time period.

Rating

Ratings can be enabled on Documents and Media. Whether the rating is a simple Thumb Up/Thumbs Down, or a Stars rating (0-5 Stars), the Rating in Analytics Cloud is a calculation that lets you compare how Users evaluate the Site's Documents and Media Files.

Audience

The Audience report uses charts to present information about the audience interaction with the asset. It answers these questions:

- How many users are interacting with my content?
- Of these users, how many are known or anonymous?
- Of the known users who interacted with my content, how many belong to Analytics Cloud Segments?
- Of the users in Segments, what are the top 15 segments?

Here are the charts in the Audience report:

Unique Visitors: A donut chart that presents the percentage of total unique visitors who are known or anonymous. The center of the chart shows the total number of unique visitors. Tooltips on each chart section show the number of users for that section (e.g, the number of known or anonymous users).

Known Individuals: A donut chart that presents the percentage of known individuals who belong or do not belong to one or more Analytics Cloud Segment. The center of the chart shows the total number of known individuals. Tooltips on each chart section show the number of users for that section (e.g, the number of users belonging to one or more Segments). Click the chart title (Known Individuals) to view a list of all the known individuals who comprise the chart's data.

Known Individuals Segments: A bar chart that shows the percentage of known individuals that comprise each Segment. The chart shows a bar for each of the top five Segments, and then aggregates the remaining Segments into the last bar. A tooltip on the last bar shows the values for each of the remaining Segments.

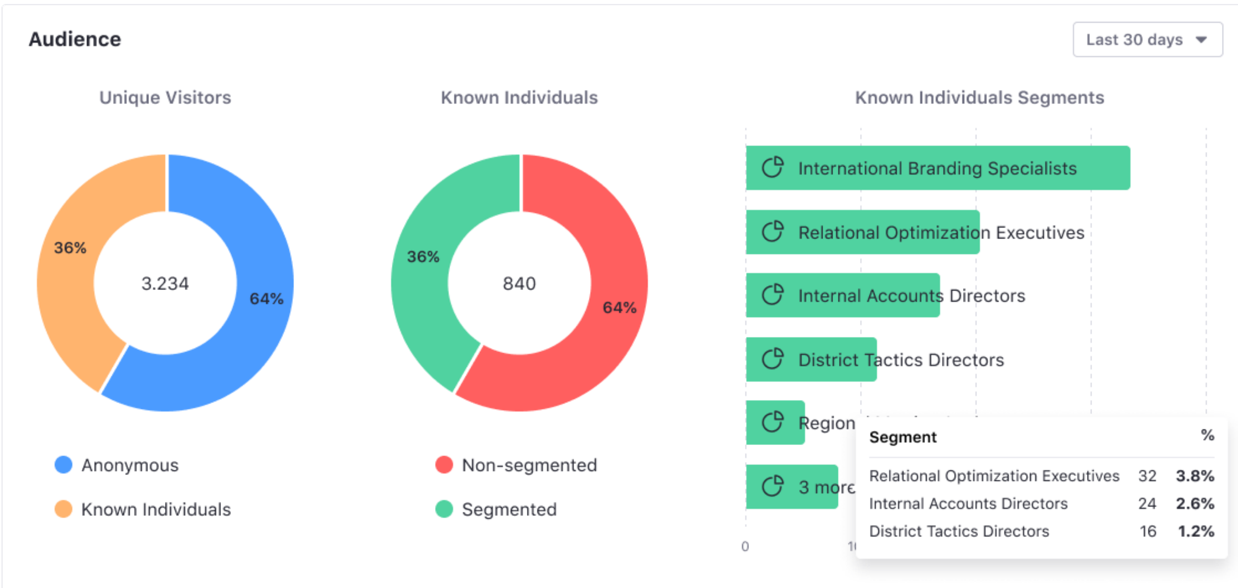

Figure 82.39: The Audience report visualizes how individuals interact with the asset.

Downloads by Location

The map in this panel shows the number of downloads by country in the selected time period.

Views by Technology

View a stacked bar graph of the Page's views by operating system (grouped by device type) in the default tab.

Click *Web Browser* to see a donut chart displaying up to the top eight web browsers over the selected time period. If applicable, remaining web browsers are aggregated in the ninth donut segment.

Asset Appears On

The Asset Appears On table includes a Page Name and URL. Like any Page data in Analytics Cloud, only Pages that were interacted with in some way by site visitors are tracked and reported. It's important to note that the link doesn't take you to the actual Page with the Asset on it. Instead, it leads to the Page Analytics view of the Page. From there you can click the URL to the actual Page.

Figure 82.40: From which location do users download the asset most frequently?

82.10 Web Content Analytics

Web Content Articles are important assets for presenting site visitors with information. Analytics Cloud gives you important insights into how site users are engaging with your content.

Visitor Behavior

The Overview tab of an Asset's detail presents several data presentations. The first chart, called Visitor Behavior, contains one important metric: total Views.

The number of views per day (or hour, if *Yesterday* or *Last 24 Hours* are selected) is plotted over the selected time period to establish a View trend line. Use this to evaluate how well the content engages visitors over time.

Audience

The Audience report uses charts to present information about the audience interaction with the asset. It answers these questions:

- How many users are interacting with my content?

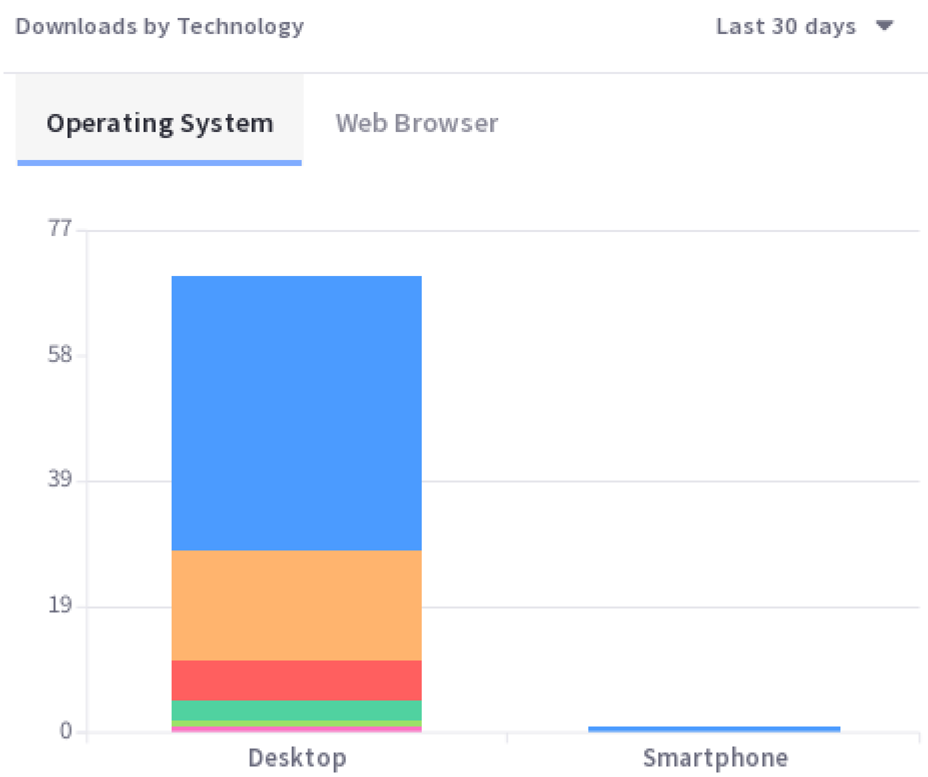

Figure 82.41: What technologies are used to interact with the Asset?

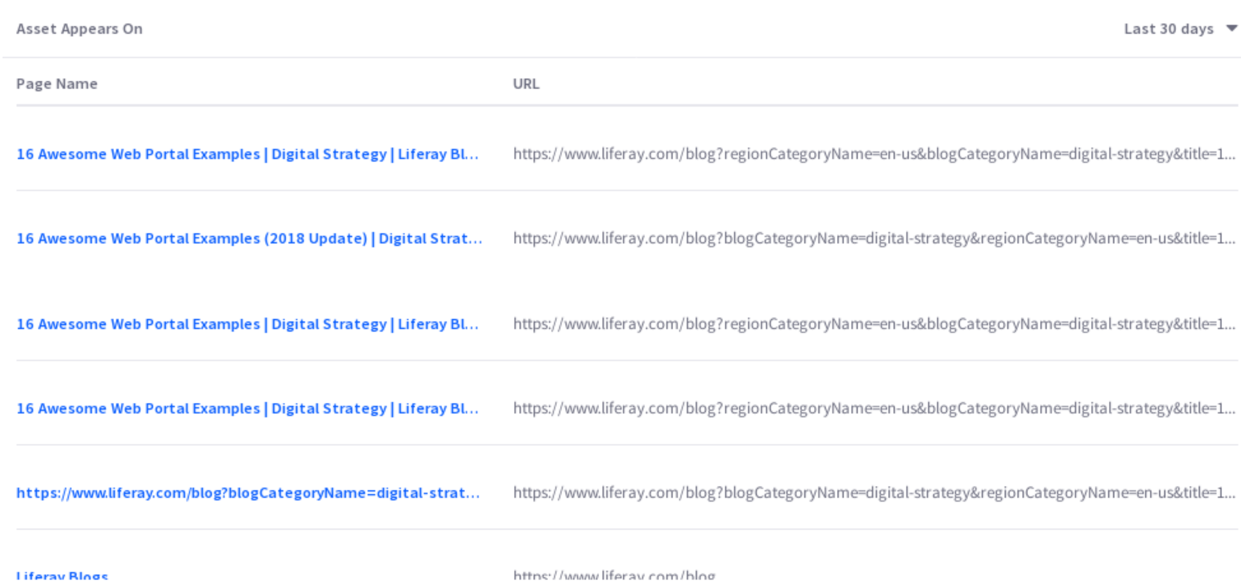

Figure 82.42: What Pages does the Asset appear on?

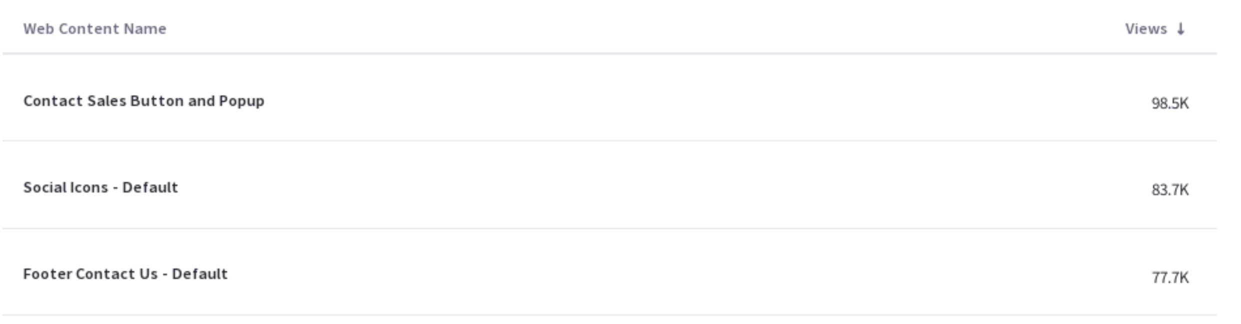

Figure 82.43: By default, Web Content is listed in descending order of Views.

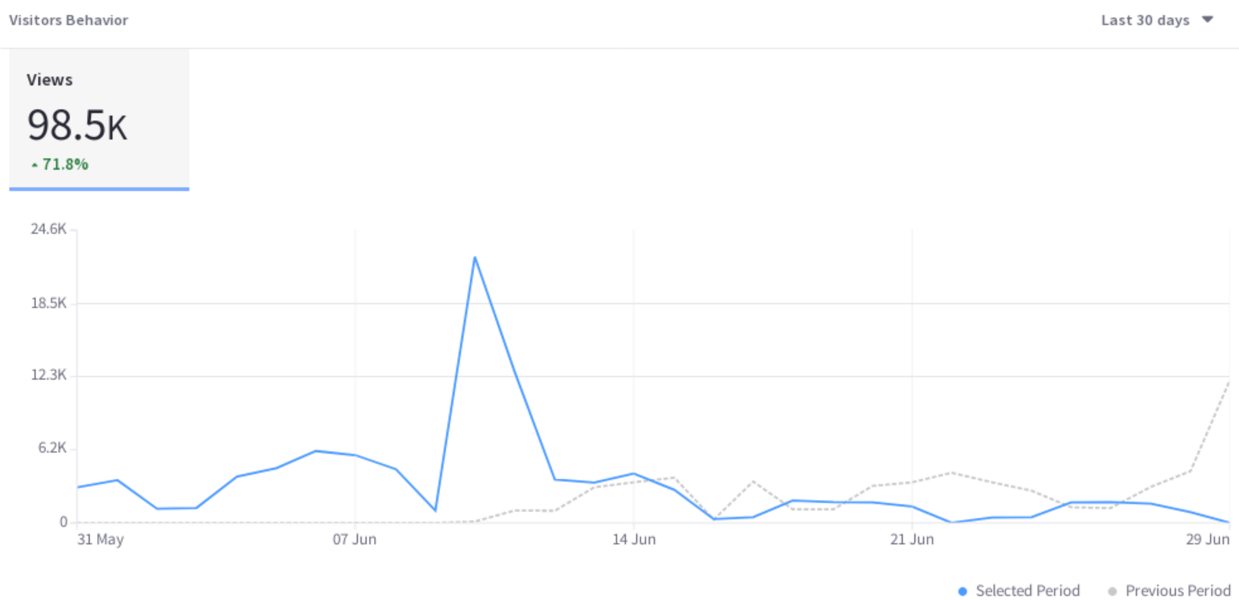

Figure 82.44: The Visitors Behavior chart contains important trend lines.

- Of these users, how many are known or anonymous?
- Of the known users who interacted with my content, how many belong to Analytics Cloud Segments?
- Of the users in Segments, what are the top 15 segments?

Here are the charts in the Audience report:

Unique Visitors: A donut chart that presents the percentage of total unique visitors who are known or anonymous. The center of the chart shows the total number of unique visitors. Tooltips on each chart section show the number of users for that section (e.g, the number of known or anonymous users).

Known Individuals: A donut chart that presents the percentage of known individuals who belong or do not belong to one or more Analytics Cloud Segment. The center of the chart shows the total number of known individuals. Tooltips on each chart section show the number of users for that section (e.g, the number of users belonging to one or more Segments). Click the chart title (Known Individuals) to view a list of all the known individuals who comprise the chart's data.

Known Individuals Segments: A bar chart that shows the percentage of known individuals that comprise each Segment. The chart shows a bar for each of the top five Segments, and then aggregates the remaining Segments into the last bar. A tooltip on the last bar shows the values for each of the remaining Segments.

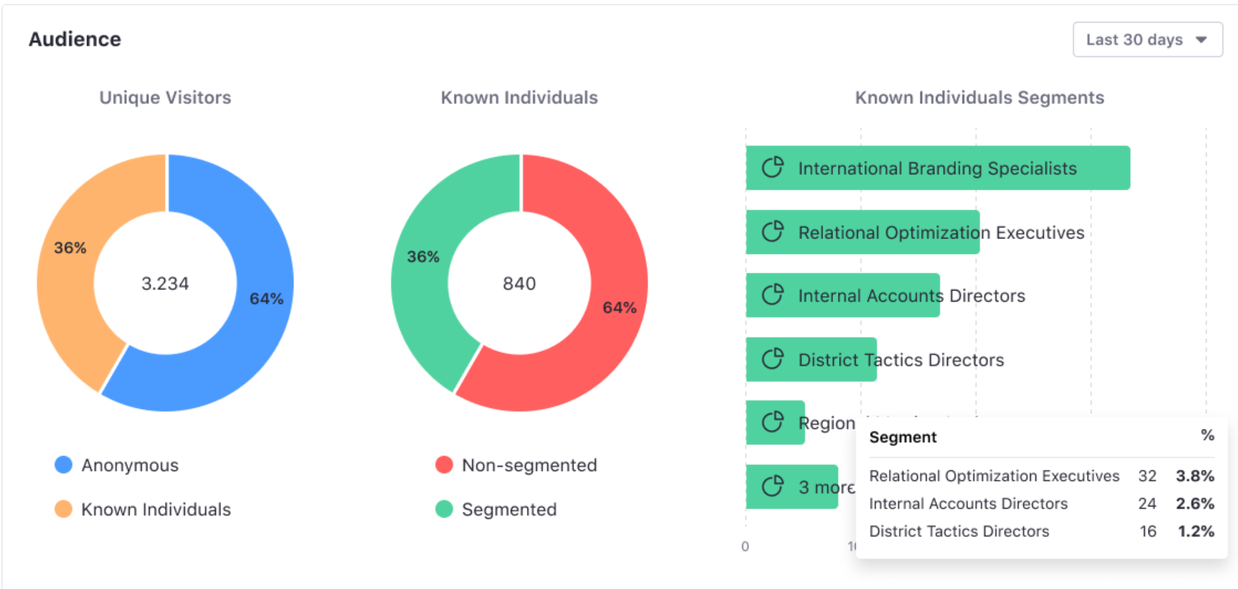

Figure 82.45: The Audience report visualizes how individuals interact with the asset.

Views by Location

The map in the Views by Location panel shows the number of views by country in the selected time period.

Views by Technology

View a stacked bar graph of the Page's views by operating system (grouped by device type) in the default tab.

Click *Web Browser* to see a donut chart displaying up to the top eight web browsers over the selected time period. If applicable, remaining web browsers are aggregated in the ninth donut segment.

Asset Appears On

The Asset Appears On table includes a Page Name and URL. Like any Page data in Analytics Cloud, only Pages that were interacted with in some way by site visitors are tracked and reported. It's important to note that the link doesn't take you to the actual Page with the Asset on it. Instead, it leads to the Page Analytics view of the Page. From there you can click the URL to the actual Page.

Figure 82.46: From which location do users interact with the Asset most frequently?

82.11 A/B Testing Analytics

A/B Testing evaluates the effectiveness of Content Pages by testing multiple versions/layouts of the Page at once. This is done by creating Page Variants of the original Page, testing the Page with a goal (e.g., clicks), and publishing the most effective Variant. You can learn more about creating an A/B test and configuring it for a Content Page in Liferay DXP's A/B Testing documentation.

All results from an A/B test running in Liferay DXP are tracked by Analytics Cloud. An A/B test is synced with Analytics Cloud once it's created. From there, you can manage the A/B test from Analytics Cloud. To view all drafted, running, terminated, and completed A/B tests, go to the *Tests* menu from the left column.

For a drafted A/B test, you can manage its

- *Target*: the Experience and User Segment.
- *Metric*: the goal to track (e.g., Bounce Rate or Click).
- *Variants*: the Page Variants for users to interact with.
- *Traffic Split*: the percentage of visitors that are randomly split between the Variants when visiting the Page.
- *Confidence Level*: the accuracy of the test results.

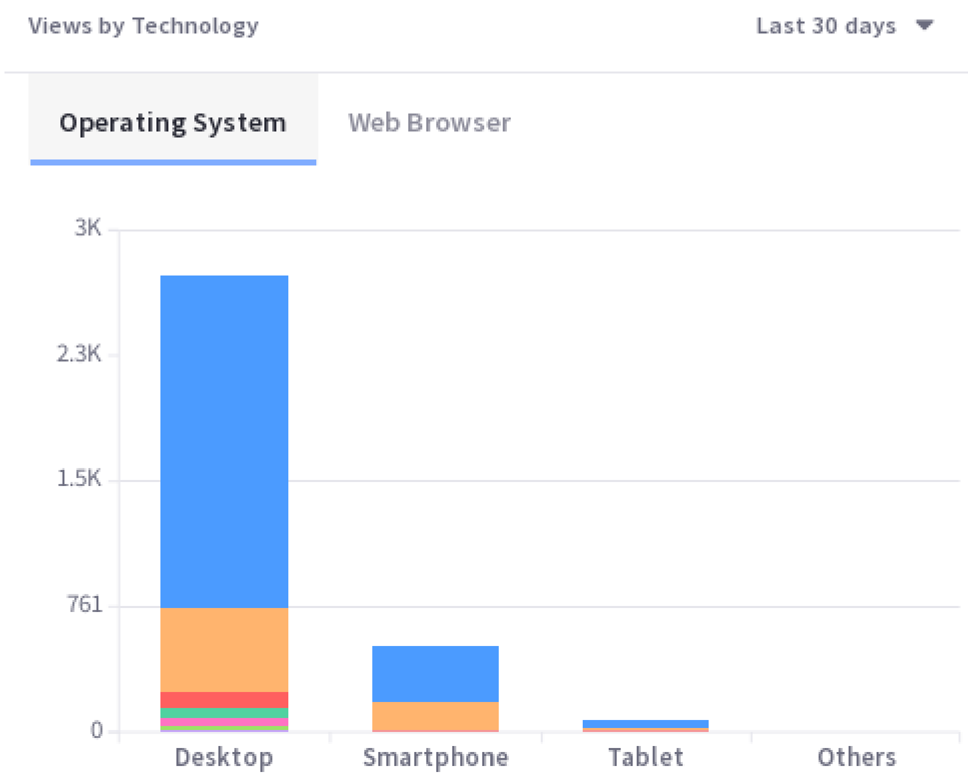

Figure 82.47: What technologies are used to interact with the Asset?

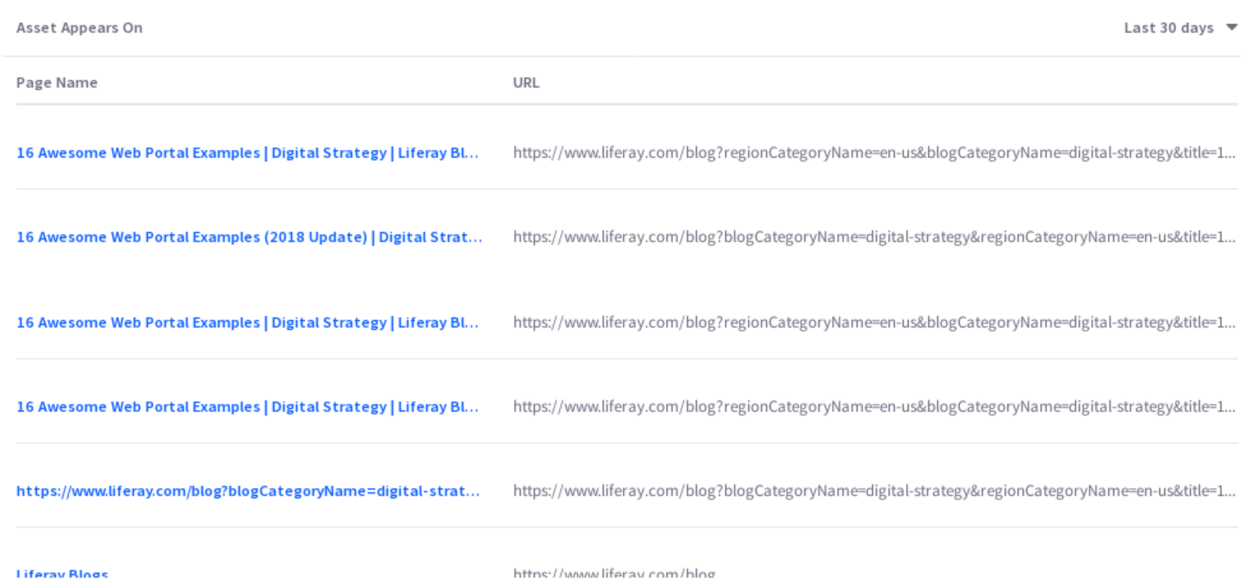
Asset Appears On

Last 30 days

Page Name	URL
16 Awesome Web Portal Examples | Digital Strategy | Liferay Bl...	https://www.liferay.com/blog?regionCategoryName=en-us&blogCategoryName=digital-strategy&title=1...
16 Awesome Web Portal Examples (2018 Update) | Digital Strat...	https://www.liferay.com/blog?blogCategoryName=digital-strategy&regionCategoryName=en-us&title=1...
16 Awesome Web Portal Examples | Digital Strategy | Liferay Bl...	https://www.liferay.com/blog?regionCategoryName=en-us&blogCategoryName=digital-strategy&title=1...
16 Awesome Web Portal Examples | Digital Strategy | Liferay Bl...	https://www.liferay.com/blog?regionCategoryName=en-us&blogCategoryName=digital-strategy&title=1...
https://www.liferay.com/blog?blogCategoryName=digital-strat...	https://www.liferay.com/blog?blogCategoryName=digital-strategy&regionCategoryName=en-us&title=1...
Liferay Blogs	https://www.liferay.com/blog

Figure 82.48: What Pages does the Asset appear on?

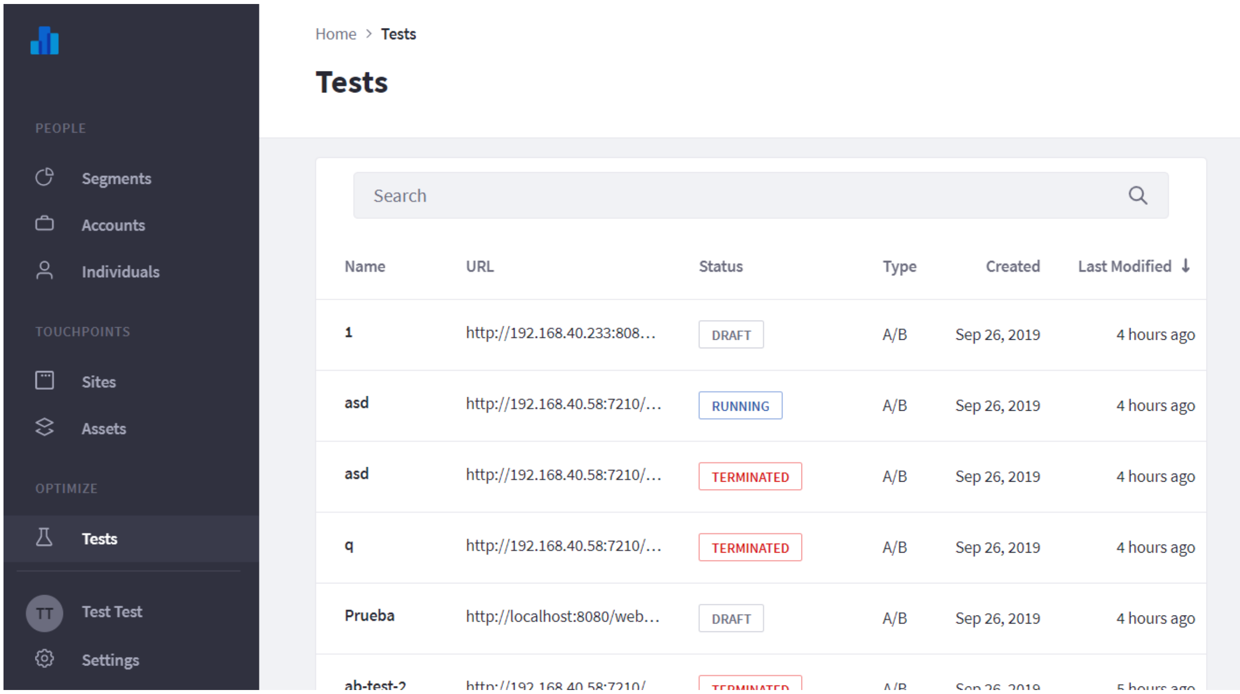

Figure 82.49: A complete history of your A/B tests are available in Analytics Cloud.

See Liferay DXP's A/B Testing documentation for more information on an A/B test's setup.

Once your A/B test is running, Analytics Cloud offers several reports to keep you up-to-date on your A/B test's progress:

- *Summary*
- *Variant Report*
- *Test Sessions*

You'll learn about these next.

Summary

The Summary panel gives you an overview of your test. It provides you with information like

- percent completion
- running time (in days)
- total visitor sessions

It also gives you a quick glance at your test metric and the best current performing Variant.

Variant Report

The Variant Report panel provides a detailed breakdown of each Variant and how well they're performing.

Below are the metrics reported for each variant:

Median: the middle number in the set of sample values. This estimates a typical user's behavior.

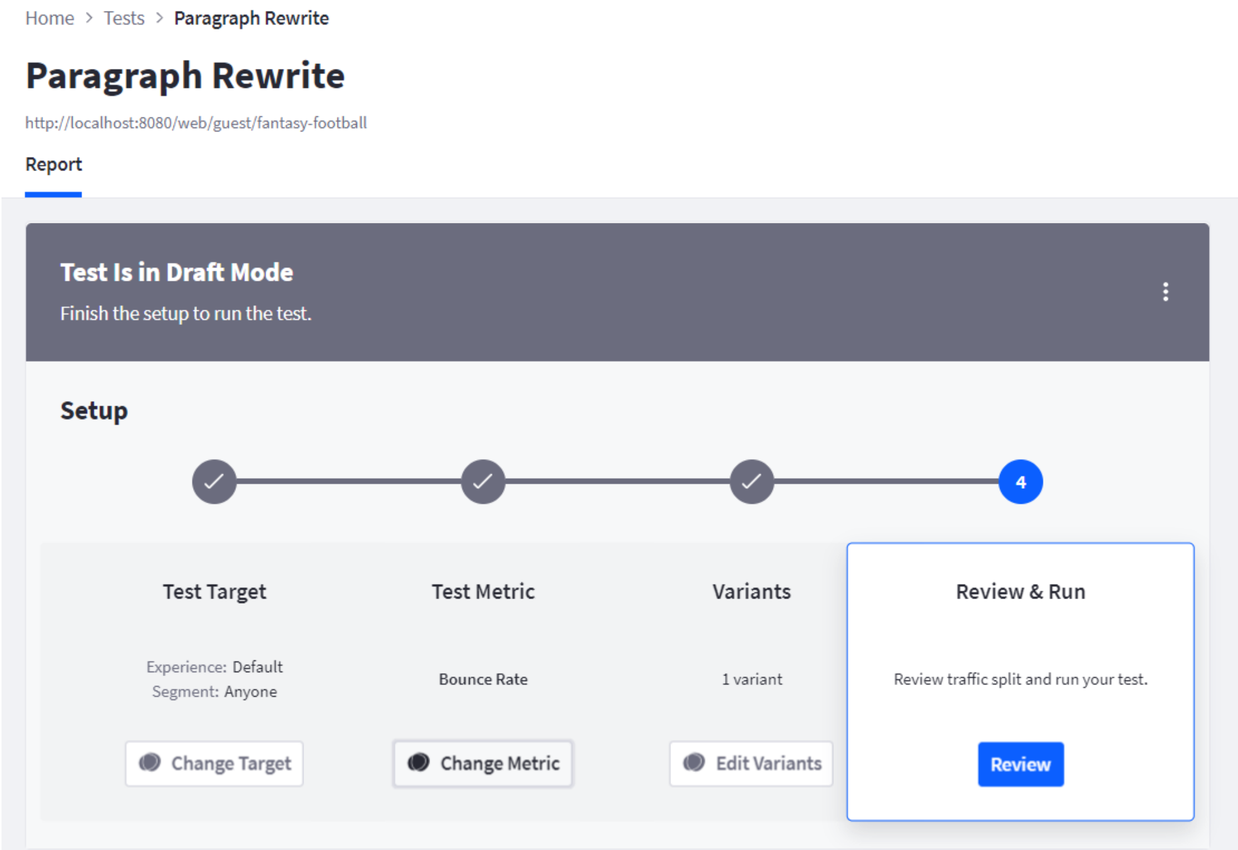

Figure 82.50: Analytics Cloud offers a visual checklist for an A/B test's setup.

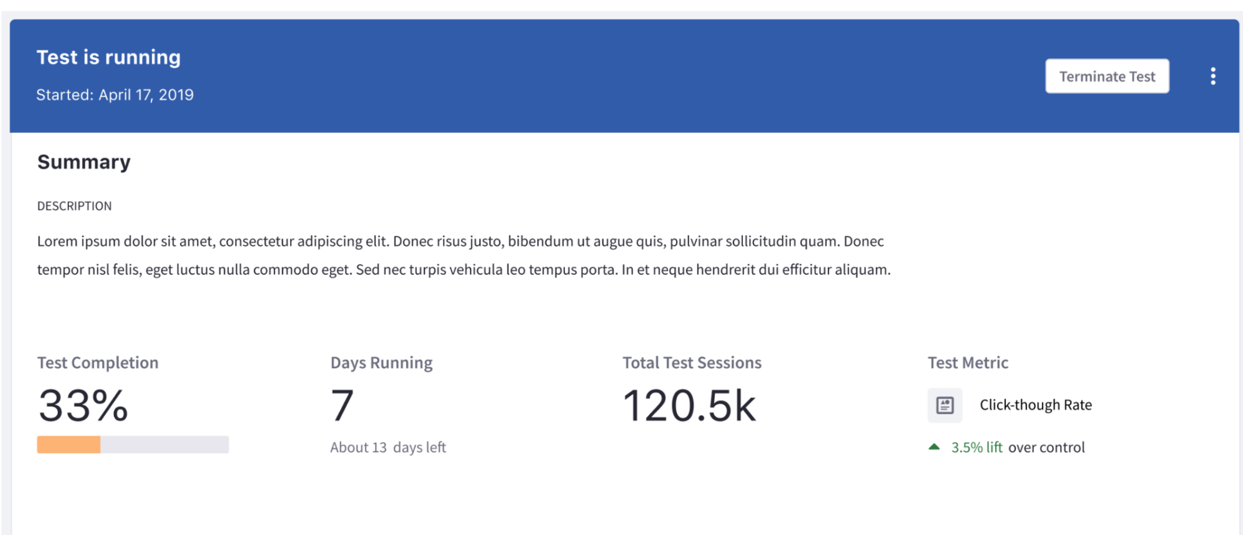

Figure 82.51: The Summary panel gives you a quick way to assess the status of your A/B test.

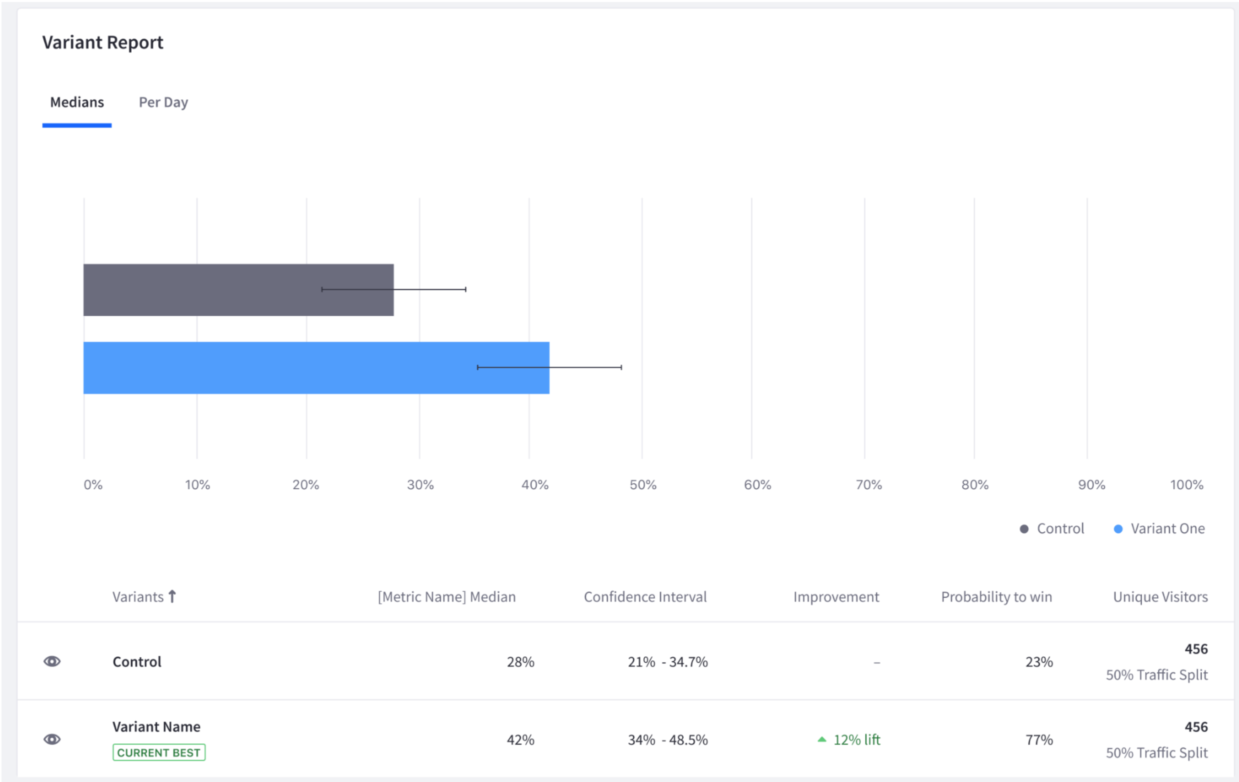

Figure 82.52: Variants are tracked using multiple metrics.

Confidence Interval: the range of values expected to contain the true mean of the population. For example, a 95% confidence interval is a range of values that the system is 95% sure contains the true mean. This gives the range of possible values that seem plausible for the measured goal.

Improvement: the relative improvement from the control group. This metric may also be known as *Lift*. For example, assume the Control Page has a 15% retention rate. The improvement calculation would be `((16 - 15) / 15) = ~6.67%` improvement.

The lets you know the impact of a change. If there is only a small improvement, it may not be worth implementing that change.

Probability to Win: predicts the likelihood that the Variant will beat out all other participating Variants. This lets you see how multiple metrics compare to each other. For example, consider a horse racing event: each horse has a generated chance to win that is posted before a race (i.e., odds of winning), calculated by simulating the race thousands of times. This same method is used for your Variants to calculate their probability of winning the A/B test.

Unique Visitors: the number of visitors contributing to the Variant. A visitor randomly assigned a Variant always sees the same Variant until the test is finished.

Besides knowing how much traffic is hitting a page, this metric also helps determine if there is an issue with how the A/B test is configured. For example, there could be too much traffic going to one Variant (typically caused by a Segment misconfiguration).

Test Sessions

The Test Sessions panel provides statistics showing how many sessions view your test impressions per day over time. This helps you validate that your audiences are being directed to your A/B test impressions. It also portrays how your test affects the traffic to your page compared to before.

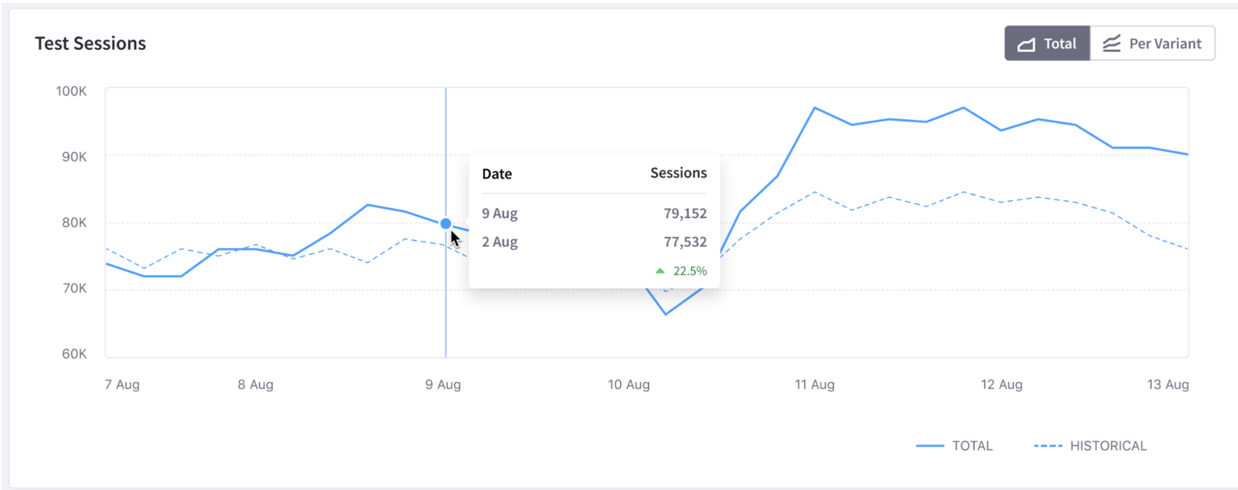

Figure 82.53: This graph gives you an accurate depiction of your page traffic.

Next, you'll learn about an A/B test's statuses.

Test Status

An A/B test is always characterized with a status after it starts. These include

- *Test is Running*
- *Winner Declared*
- *No Clear Winner*

You'll explore each status next.

Test is Running

This means that your test is still running and needs a larger sample size before declaring a winner. You can still see which Variant is your current best; however, the desired confidence level has not been met.

When a test is running, you can terminate it by selecting *Terminate* from the Summary bar.

Winner Declared

Once your A/B test successfully finishes, a Variant is declared a winner. At this state, you can perform the following actions:

- publish the winning Variant as your default experience.
- complete the test without publishing any Variants.

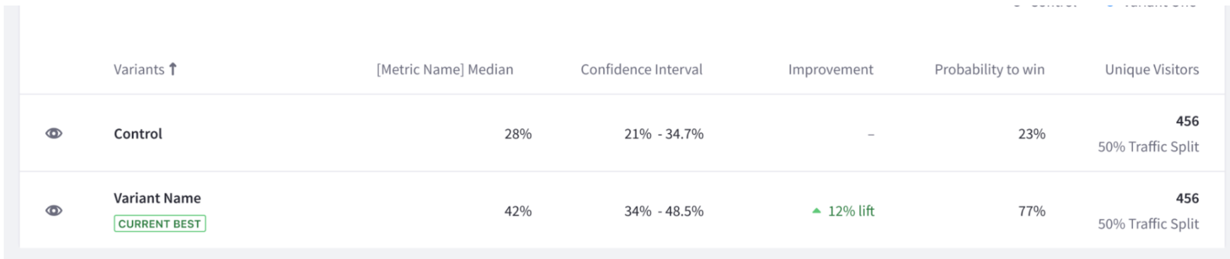

Figure 82.54: The leading Variant is clearly labeled as being the current best.

Figure 82.55: Terminating an A/B test allows you to delete the test, if desired.

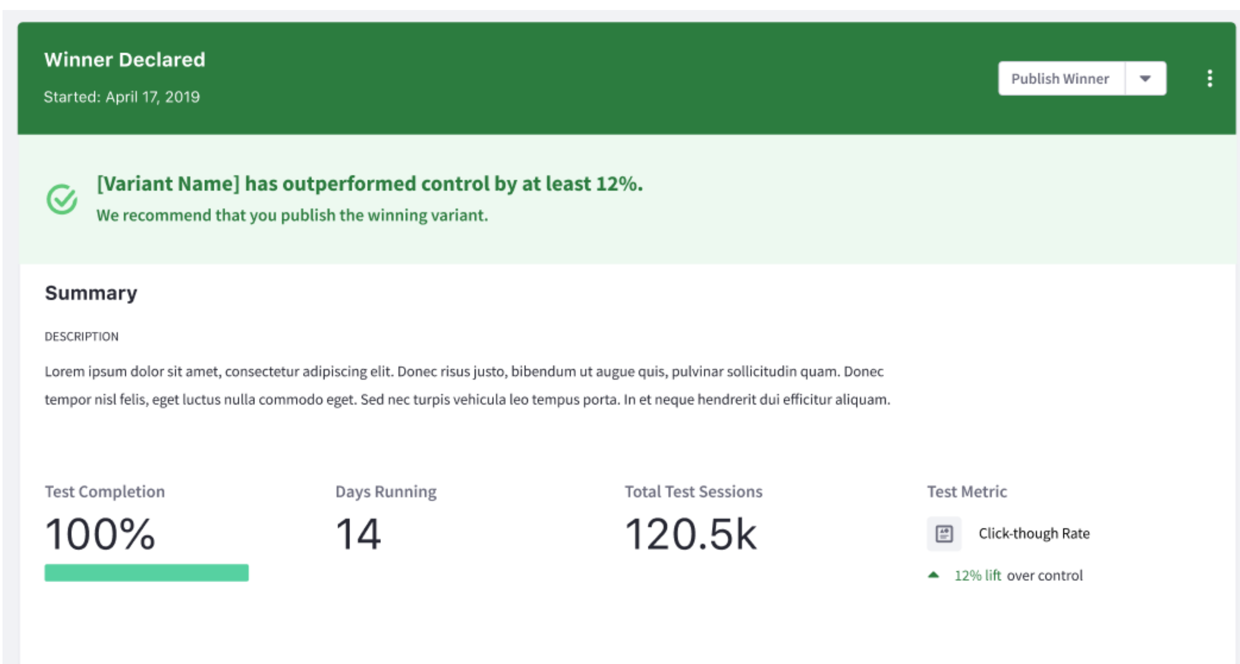

Figure 82.56: Click *Publish Winner* to publish the winning Variant.

No Clear Winner

Sometimes, Analytics Cloud cannot determine a winner because no Variant has outperformed significantly over the Control Page. In this case, you can complete the test without publishing anything. The control stays the default experience.

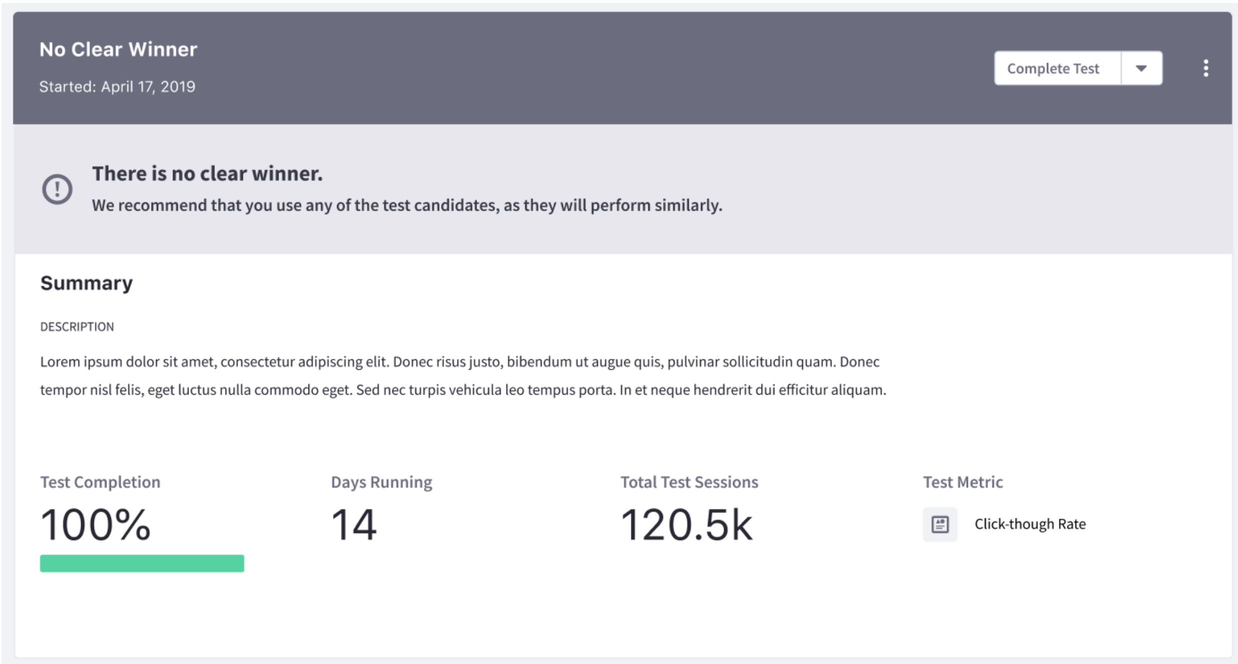

Figure 82.57: When the required confidence level is not met during the time duration, there is no winning Variant.

By viewing the generated analytics for your A/B tests, you're constantly informed on how they're progressing. With the provided data, you can confidently choose the best Experience for your Site's users.

82.12 Content Recommendation API

The Content Recommendation API in Liferay Analytics Cloud suggests content based on user interactions with content in a Liferay DXP instance. This is possible because once Liferay DXP is connected to Analytics Cloud, it sends interaction events each time a user visits a page. Each of those events contains information about the content the user consumes. Analytics Cloud collects and processes this information.

The Content Recommendation API contains services that enable the following:

- Discover similar tags (content) based on the current tags the user is browsing.
- Discover tags based on the user's interest over time.

Discover Similar Tags

This service returns a list of tags (terms) similar to those passed as parameters. The list of similar terms is sorted by weight and returned as JSON.

Here's this API's endpoint:

GET: `{url}/api/1.0/interests/terms/related{?page,size,terms}`

Here are the parameters:

`int page (defaultValue = 0)`: The page of results. For example, the default value `0` specifies the first page of results.

`int size (defaultValue = 5)`: The number of results to include on each page. The default value `5` specifies five results on each page.

`List<String> terms`: The tags to use for determining the list of similar tags. This is the only required parameter.

`double termWeightThreshold (defaultValue = 0.01)`: The relevance level (weight) for determining related terms. The default value of `0.01` returns all tags with a weight above 1%.

Together, the `page` and `size` parameters control the number of similar terms to include in the response. Note that this API can return up to 100 terms. If the number of terms exceeds that limit, the API returns an error.

For example, here's a JSON response that contains related terms:


```
{
  "_embedded": {
    "interest-terms": [
      "jquery",
      "html",
      "sql",
      "mysql",
      "java"
    ]
  },
  "page": {
    "number": 0,
    "size": 5,
    "totalPages": 7,
    "totalElements": 35
  }
}
```


Discover Tags Based on Interest

This service returns a list of tags (terms) that are relevant to a specific user. You must pass that user's ID as a parameter. The list of similar terms is sorted by weight and returned as JSON.

Here's this API's endpoint:

GET: `{url}/api/1.0/interests/terms/{userId}`

There are also three optional parameters you can use. You can use these parameters to fine-tune the terms that the API returns. Note that topics are groups of terms:

`int termsPerTopic (defaultValue = 3)`: The number of terms to consider per topic/subject.

`double termWeightThreshold (defaultValue = 0.01)`: The relevance level for determining terms of interest. The default value of `0.01` returns all terms with a weight above 1%.

`int topicsLength (defaultValue = 3)`: The number of topics to consider.

Decreasing `termsPerTopic` and increasing the `topicsLength` might lead to an increment of the subject variation (terms from different topics being recommended to users).

Here's an example request that contains only a user ID:

```
{url}/api/1.0/interests/terms/953be104-5540-abf8-59b8-55f895200acc
```

And here's an example response in JSON:

```
{
  "_embedded": {
    "interest-topics": [
      {
        "terms": [
          {
            "weight": 0.0945945945945946,
            "keyword": "javascript"
          },
          {
            "weight": 0.08648648648648649,
            "keyword": "jquery"
          },
          {
            "weight": 0.07027027027027027,
            "keyword": "html"
          }
        ],
        "weight": 0.08653350323695352,
        "id": 7
      },
      {
        "terms": [
          {
            "weight": 0.1322314049586777,
            "keyword": "php"
          },
          {
            "weight": 0.06060606060606061,
            "keyword": "sql"
          },
          {
            "weight": 0.05509641873278237,
            "keyword": "mysql"
          }
        ],
        "weight": 0.08027610626914822,
        "id": 1
      },
      {
        "terms": [
          {
            "weight": 0.15204678362573099,
            "keyword": "java"
          },
          {
            "weight": 0.10526315789473684,
            "keyword": "android"
          },
          {
            "weight": 0.023391812865497075,
            "keyword": "multithreading"
          }
        ],
        "weight": 0.07511374008317741,
        "id": 9
      }
    ]
  }
}
```

Chapter 83

Managing Your Analytics Project

Project administrators connect data sources to the project, manage project users, and track project usage against the current plan limits.

The Getting Started articles demonstrate these processes:

- Adding Liferay DXP and CSV data sources.
- Managing Data Sources

The following articles demonstrate these administrative responsibilities:

- Managing Users: As you're starting your project and new people join your team, you'll want to invite teammates to participate in your project. This article shows you how to invite users and assign them appropriate permission Roles.
- Tracking Usage: As you sync contact and web analytics data from data sources, you must keep an eye on your project's usage compared to the Liferay Analytics Cloud plan you purchased. The Usage page shows your current usage and describes available plan upgrades and add-ons to address your analytics needs.

Next Managing Users shows you how to get your team on-board your project.

83.1 Managing Users

If you're the project *Owner* (Liferay invited you by email) or have been assigned the *Admin* permission, you can invite and manage users. Here's how to bring up the User Management page:

1. Click on *Settings* in the navigation panel. The default Settings page appears.
2. Click on *User Management* in the navigation area. The *User Management* page appears, listing all the project's users.

From the *User Management* page, an Admin user (or Owner) can invite users, edit their permissions, and delete them.

If your project is new, you'll want to invite your teammates to the project.

Inviting Users

It's exciting to give and receive invitations. Carry that excitement with you as you invite your team's analysts, data integrators, and marketers to your Analytics Cloud project.

1. In the *User Management* page, click on the *Invite Users* button. The *Invite Users* dialog appears.
2. Enter the email addresses (separated by a space or comma) of the people to invite to the project.
3. Click the *Send* button. The invitation is sent to your lucky teammates.

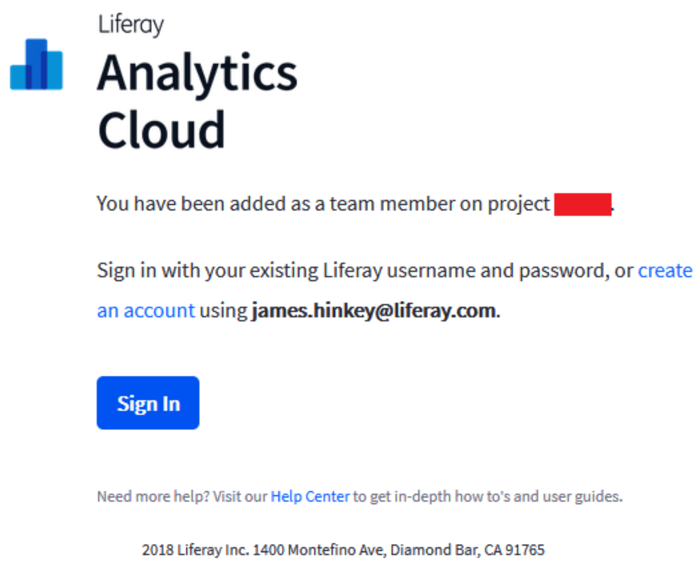

Figure 83.1: The invitation email looks like this, except your project number shows in place of the red rectangle.

You can set the user's permissions when their invitation acceptance is pending.

Managing Permissions

Analytics Cloud projects have these permission Roles:

Owner: Can do all things, but cannot be deleted.

Admin: Can do all things, including manage other Admin users. Admin users can be deleted.

Member: Can view all profiles and analytics, but cannot edit anything under Settings. They cannot manage users, or add or remove data sources.

Here's how to change a user's permission Role:

1. Click the *Edit* button in the user's list entry.
2. In the *Permission* column, select the user's current permission. A permission Role selector appears.
3. Select the permission Role for the user.

The user is assigned the selected Role. That's all there is to it!

83.2 Tracking Usage

Analytics Cloud plans are limited by the total amount of Individuals and Pages Views synced from data sources. When either limit is exceeded, you must either upgrade your plan or purchase add-ons to accommodate the overage.

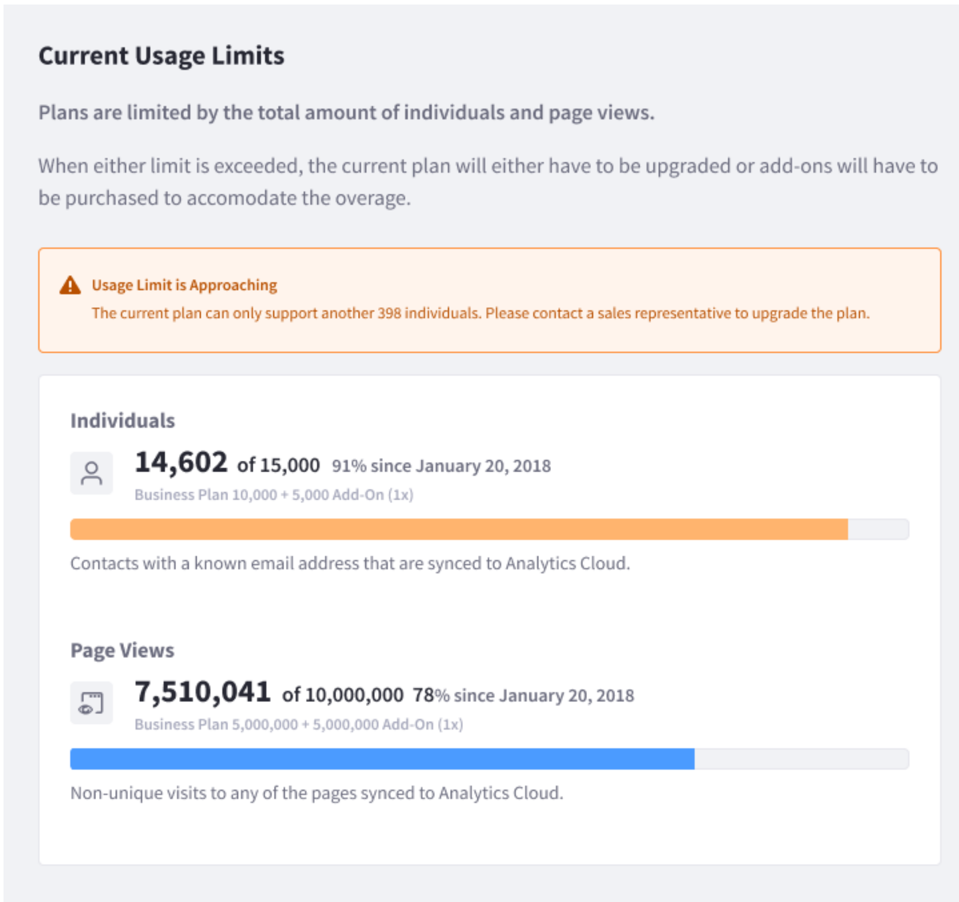

Figure 83.2: The Usage page reports when you're approaching or exceeding your plan limits.

The plans and add-ons are described on the right side of the page.

Plans

The *Plans* section describes each plan's limits and cost. The higher tier plans cost more initially but offer more flexibility through higher limits and less expensive add-ons. Your current plan is highlighted. You can upgrade from lower tier to higher tier plans.

Add-ons

Add-ons increase limits on Individuals or Page Views, without requiring you to upgrade to a new plan. Purchasing add-ons gives you commitment-free analytics capacity boosts.

Contact your sales representative for further guidance on getting a plan that fits you.

83.3 Managing Interest Topics

Analytics Cloud's recommendation API suggests topics of interest based on user interactions with content in Liferay DXP. Analytics Cloud administrators can exclude certain keywords from the recommendation API, therefore preventing Analytics Cloud from suggesting content based on those keywords. This is useful if the administrator considers a given topic too broad, too narrow, offensive, or otherwise not useful for suggesting content.

Follow these steps to add a keyword to the block list:

1. In Analytics Cloud, select *Settings* → *Interest Topics*. A table lists any existing keywords in the block list.
2. Click *Add Keyword* and enter one or more keywords in a comma-delimited list.
3. Click *Send* when you're finished. Your keywords then appear in the table.

To delete a keyword from the block list, click that keyword's trash icon ( ). To delete multiple keywords at once, select the checkbox for each and click the trash icon that appears at the top of the table. You can select all keywords via the checkbox at the top of the table.

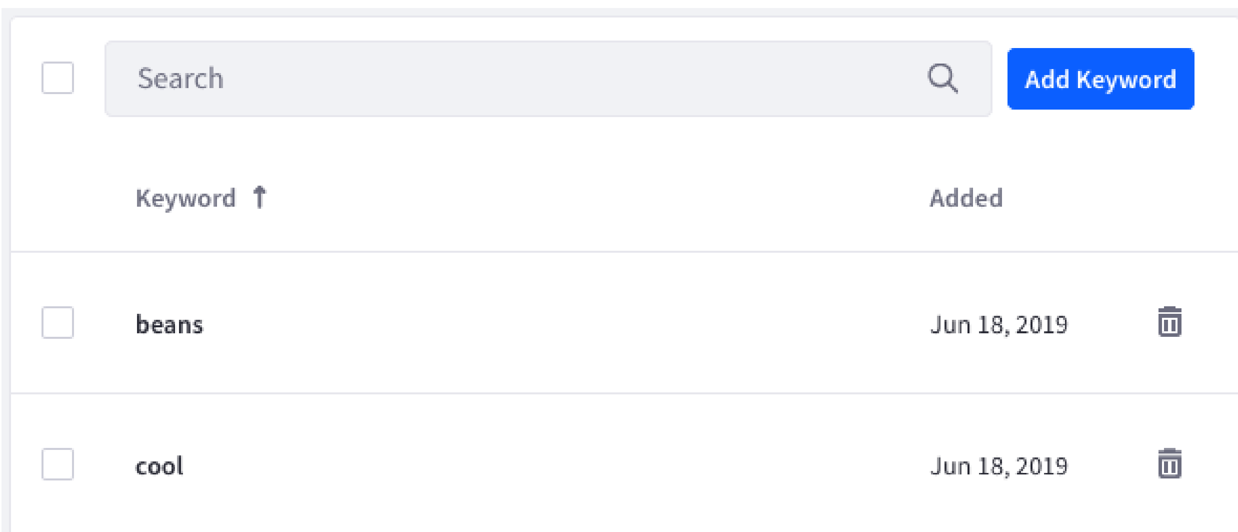

Figure 83.3: Cool beans aren't allowed.

Chapter 84

Troubleshooting Liferay DXP Data Sources

Misconfigured environments or data sources can prevent or disrupt access to Liferay DXP (DXP) data sources. Here's how to troubleshoot DXP data source issues.

84.1 Retry Authorization

Error Message: `Unknown error. Please retry authorization.`

DXP data source access requires that your DXP instance be publicly available and that your Analytics Cloud instance be registered with the DXP instance as an OAuth application.

Resolution:

1. Follow the steps for adding a Liferay DXP data source.

2. Register Analytics Cloud with your DXP instance.

84.2 Unsupported Version

Error Message: `Unsupported version. This method of connection does not support the data source Liferay version. Make sure you are connecting to Liferay 7.0/7.1 instance or try a different method of connection.`

Resolution:

1. Make sure to connect with a Liferay DXP 7.0 or 7.1 instance.

2. Follow the steps for adding a Liferay DXP data source.

3. If the error persists, make sure JSON web services are enabled on your DXP instance. They're enabled by default. If you disabled them using a portal property setting `json.web.service.enabled=false` (e.g., set in a `portal-ext.properties` file), delete the setting or set the property value to `true`.

Invalid Credentials; the Authorization Expired

Error Message: `Invalid Credentials. The authorization for this data source has expired. Please reauthorize the token in the OAuth tab.`

This message appears in the log:

`WARN [http-nio-8080-exec-14][AbstractOAuthService:88] Unsecure HTTP, Transport Layer Security is recommended`

Connection to a DXP data source requires that the DXP instance's web server protocol be forwarded.

Resolution:

1. Follow steps for adding a DXP data source, paying particular attention to register Analytics Cloud with your DXP instance.

2. If the issue persists and the web server protocol is forwarded, set these portal properties in a `portal-ext.properties` file in your DXP instance.

```
web.server.forwarded.protocol.enabled=true
redirect.url.security.mode=domain
redirect.url.domains.allowed=
```

Part III

Deploying Liferay DXP

CHAPTER 85

DEPLOYING LIFERAY DXP

Liferay DXP is one of the most flexible applications on the market today with respect to database and application server support. Liferay DXP supports a wide variety of databases and application servers, freeing you to use the ones you think are best. Liferay DXP also scales very well. You can install Liferay DXP on a shared hosting account, on a multi-node cluster running a commercial application server, or on anything in between. In fact, Liferay DXP is used successfully in all of these scenarios every day.

You'll find that because Liferay DXP is extremely flexible in its deployment options, it is also easy to install. If you already have an application server, you can use your application server's deployment tools to install Liferay DXP. If you don't already have an application server, Liferay provides several application server bundles from which to choose. These are pre-configured application servers with Liferay DXP already installed on them. They're easy to install and with a small amount of configuration can be made into production-ready systems.

Read on to learn how you can prepare to install Liferay DXP.

85.1 Preparing for Install

Installing Liferay DXP is easy. But before you begin, you should answer a few questions.

- Which version of Liferay will you install?
- On which application server will you install Liferay DXP?
- Can you install a Liferay DXP bundle or do you need to install Liferay DXP on an existing application server?
- Which database will you use with Liferay DXP?
- How do you plan to store your data?
- Can your network support Liferay DXP?

Next, you'll answer these questions and learn the basic steps for installing Liferay DXP.

Obtaining Liferay DXP

Anyone can download Liferay DXP from liferay.com. Click Product → *Downloads*, and you'll be able to download either the open source version of Liferay DXP or a trial of the commercial version in several different formats. These include our convenient bundles as well as `.war` files for installing Liferay DXP on your application server of choice.

Liferay enterprise subscribers can download Liferay DXP from the Help Center.

So what is a Liferay DXP bundle anyway? A Liferay DXP bundle is an application server with Liferay DXP pre-installed. Bundles after Liferay DXP 7.1 GA1 and Liferay Portal CE 7.1 GA1 are released in 7-Zip (`.7z`) and gzipped tar (`.tar.gz`) compressed file archive formats. Bundles are the easiest way to install Liferay DXP. Liferay DXP is bundled with a number of application servers; choose the one that best fits your needs. If you don't currently have an application server preference, consider starting with the Tomcat bundle. Tomcat is one of the most lightweight and straightforward bundles to configure. If you have an open source application server preference, choose the server you prefer from the available Liferay DXP bundles. Liferay DXP requires a Java JDK 8 or 11.

Note: Please see the compatibility matrix for information on supported JDKs, databases, and environments.

Please note that Liferay DXP is not able to provide application server bundles for proprietary application servers such as WebLogic or WebSphere, because the licenses for these servers don't allow for redistribution. Liferay DXP's commercial offering, however, runs just as well on these application servers as it does on the others. You must follow our manual installation procedure using a `.war` file to install Liferay DXP on proprietary application servers.

Once you have Liferay DXP, you can plan out your installation. First, prepare your database. Second, install Liferay DXP. Third, configure your network. Fourth, configure search. You can install Liferay DXP either by using a bundle or by installing it manually on your existing application server. Next, we'll go over the steps it takes to install Liferay DXP.

Liferay DXP Installation Steps

Before you begin installing Liferay DXP, you should review these basic installation steps:

1. Choose a database server to use with Liferay DXP and create a new database. Determine whether you want Liferay DXP to manage your database connection or your application server to manage your database connection. We recommend that you let Liferay DXP manage your database connection. Liferay DXP can connect with several open source or enterprise level document repositories.

2. Gather mail credentials for sending email notifications to users. Determine whether you want Liferay DXP to manage your mail session or your application server to manage your mail session. Liferay DXP provides a built-in mail session but also supports a JNDI mail session. We recommend that you let Liferay DXP manage your mail session.

3. Install either a Liferay DXP bundle or install Liferay DXP on an existing application server (further instructions below).

4. Choose IPv4 or IPv6. Determine which address format is best for your network (further instructions below).

5. Determine how you'll configure Elasticsearch. Liferay DXP's default embedded configuration is not supported for production use, so you must install Elasticsearch separately, either on the same infrastructure or on its own.

6. Configure ports (optional). Liferay's application server (e.g., Tomcat or Wildfly) uses certain ports for purposes like handling incoming HTTP requests, HTTPS requests, or AJP requests, etc. If you start your application server in debug mode, there's a port listening for a debugger to connect. If desired, you can configure these ports. Please refer to your application server's documentation for information on its default ports and how to configure them.

   Liferay also provides access to its OSGi framework through a configurable port:

   `module.framework.properties.osgi.console=localhost:11311`

   You can override this default property by copying the line above to your `LIFERAY_HOME/portal-ext.properties` file and adjusting the port number.

We'll go through the steps in order, so first we'll look at the Liferay DXP database.

Step 1: Choose a Database Server and Create a New Database

The recommended way of setting up your Liferay DXP database is also the simplest. Liferay DXP takes care of just about everything. You only need to take two simple steps:

1. Create a blank database encoded with the character set UTF-8. Liferay DXP is a multilingual application and needs UTF-8 encoding to display all of its supported character sets.

```
**Note:** If you plan to migrate from one database vendor to another,
[configure the database to use the default query result order you expect for entities Liferay DXP lists](/docs/7-0/tutorials/-
/knowledge_base/t/sort-order-changed-with-a-different-database).
```

```
**Note:** If you use Sybase, configure the database to allow null values by
default.
```

---

2. Create a database user for accessing this database. Grant this database user all rights, including the rights to create and drop tables, to the blank Liferay DXP database.

```
**Important:** Liferay requires reading from and writing to the database.
The Liferay database user must therefore have permissions to read and
write data.
```

Liferay DXP uses this database user's credentials to connect to the Liferay DXP database either directly or through its application server. During its initial startup, Liferay DXP creates the tables it needs in the database you just created. It does this automatically, complete with indexes.

This is the recommended way to set up Liferay DXP. It enables Liferay DXP to maintain its database automatically during upgrades or when various Liferay DXP plugins that create database tables of their own are installed. This method is by far the best way to set up your Liferay DXP database.

Warning: If you're using an Oracle database, use the `ojdbc8.jar` driver library with at least Oracle 12.2.0.1.0 JDBC 4.2 versioning because data truncation issues have been detected reading data from CLOB columns.

If you choose to set up Liferay DXP's database with the recommended permissions described in this section, you can skip to the next section.

Warning: The instructions below are not ideal for Liferay DXP installations. This procedure is documented here so that enterprises with more restrictive standards can install Liferay DXP with stricter (but sub-optimal) database settings. If it's at all possible, we recommend that you use the method described in the previous section instead of the procedure outlined below.

Even though Liferay DXP can create its database automatically, some enterprises prefer *not* to allow the database user configured in an application server to have the database permissions necessary for Liferay DXP and its plugins to maintain their tables. For these organizations, Select, Insert, Update and Delete are the only allowed permissions. Thus, in this section we explain how to set up the database manually. If your organization *is* willing to grant the Liferay DXP database user the permissions to create and drop tables in the database–and this is the recommended configuration–then simply use the recommended configuration described in the previous section.

1. Create a new, blank, database for Liferay DXP.

2. Grant full rights to do anything to the Liferay DXP database to the Liferay DXP database user.

3. Install Liferay DXP and start it so that it automatically populates its database.

4. Once the database has been populated with the Liferay DXP tables, remove the permissions for creating and dropping tables from the Liferay DXP database user.

There are some caveats to running Liferay DXP like this. Many Liferay DXP plugins create new tables when they're deployed. Additionally, Liferay DXP has an automatic database upgrade function that runs when Liferay DXP is upgraded. If the Liferay DXP database user doesn't have enough rights to create/modify/drop tables in the database, you must grant those rights to the ID before you deploy one of these plugins or start upgrading Liferay DXP. Once the tables are created or the upgrade is complete, you can remove those rights until the next deploy or upgrade. Additionally, your own developers might create plugins that need to create their own tables. These are just like Liferay DXP's plugins that do the same thing, and they cannot be installed if Liferay DXP can't create database tables. If you wish to install these plugins, you will need to grant rights to create tables in the database each time before you attempt to install them.

Liferay DXP offers several configurations to store Documents and Media files by setting the `dl.store.impl=` property. Available options are Simple File System Store, Advanced File System Store, CMIS Store, DBStore, JCRStore, and Amazon S3Store. In addition, Liferay DXP can be connected to various open source and enterprise-level document repositories. All of the repositories are connected to Liferay DXP through hooks available on Liferay Marketplace (see below).

Once you have your database and document repository ready, you can install Liferay DXP on your server.

Step 2: Gather Your Mail Credentials

Liferay DXP uses a mail server to send email notifications. As part of the install, therefore, you will need to have credentials that Liferay DXP can use to connect to your mail server. Specifically, you'll need to have the following information:

- Incoming POP Server and port
- POP User Name
- POP Password
- Outgoing SMTP Server and port
- SMTP User Name
- SMTP Password
- A list of JavaMail properties that would override a default configuration

Once you've gathered this information, you're ready to move on to the next step.

Step 3: Install

The next step is to install Liferay DXP. You can do this in one of two ways: by installing Liferay DXP bundled with an application server, or by installing Liferay DXP manually on an existing application server. Each Liferay DXP installation's Liferay Home contains several folders.

By far the easiest way to get Liferay DXP installed is to use a bundle. Though bundles are pre-packaged for demo purposes, it is very easy to turn them into full, production-ready Liferay DXP instances.

Step 4: Network Configuration

Liferay DXP supports both IPv4 and IPv6 address formats. You must choose between IPv4 or IPv6. By default, Liferay DXP uses IPv4 addresses. If you are using IPv6, you will have to configure Liferay DXP. There are two simple steps.

1. In the application server's environment settings, set `-Djava.net.preferIPv4Stack=false`.

2. Create a `portal-ext.properties` file in your portal's Liferay Home folder (if one does not already exist) and set the `tunnel.servlet.hosts.allowed` property to the target hosts you want to allow (e.g., *0:0:0:0:0:0:0:1*).

Step 5: Configure Elasticsearch

Liferay DXP by default ships with an embedded version of Elasticsearch. While this configuration works well for demo purposes, it is not supported in a production installation. After you install Liferay DXP, you'll need to configure it to connect to a standalone Elasticsearch server or cluster. Depending on the size of your installation, this standalone instance of Elasticsearch can reside either on the same machine you have Liferay DXP on or a different machine. For performance purposes, it is better to install Elasticsearch on a separate machine.

Step 6: Liferay Marketplace and Portal Security

The Liferay Marketplace is an integral part of the Liferay DXP experience. The Marketplace plugin is required to be installed on Liferay DXP. The Marketplace plugin enables a host of features that extend beyond just access to the online Liferay Marketplace. Some of the key features the Marketplace plugin enables are

- Liferay Marketplace: direct access to our online Marketplace
- App Manager: ability to install, uninstall, and update apps
- Bundled Apps: easily manage apps that may come bundled with your Liferay DXP
- Developer Apps: ability to manage apps that you're developing
- License Manager: streamlined license management for your Liferay DXP and apps

The portal installation process deploys and registers the Marketplace plugin automatically. If you're installing Liferay DXP in an environment that would prevent this from happening, you'll have to perform one of several workarounds.

Now that you know where you're headed, you can install Liferay DXP. If you have decided to install Liferay DXP using a bundle, continue with the next section. If you're installing Liferay DXP manually, skip to the section for your application server of choice. Once you have Liferay DXP installed manually or via a bundle, you can move on to using Liferay DXP's Setup Wizard, configuring mail, and installing Elasticsearch.

85.2 Installing Liferay DXP

Now that you've performed the steps needed to prepare for your installation, you're ready to install Liferay DXP! Since bundles are the easiest way to complete an installation, all the installation steps below assume you're installing a Liferay DXP bundle. If you plan to install Liferay DXP manually, please refer to the article for your app server of choice, and then come back here to complete the configuration steps.

Now you're ready. You've created a blank database for Liferay DXP and have gathered the credentials you need for your mail server. The next step is to install Liferay DXP.

Liferay Home

Liferay DXP bundles contain the same folder structure regardless of application server. The top-level folder is named for the Liferay DXP release. This folder is called *Liferay Home*. This folder is usually the application server's parent folder. This is why Liferay DXP bundles place the application server inside the bundle's root folder. On a manual installation, the location of this folder varies by application server. In a bundle, it's part of the bundle. If you're doing a manual installation, please refer to the article covering that app server for its location.

Liferay Home has folders for various purposes:

- **Liferay Home**
  - **[Application Server]**: The name of this folder varies depending on the bundle you're using. This folder contains the application server in which Liferay DXP has been installed.

- data: Stores an embedded HSQL database, Liferay DXP's file repository, and Liferay DXP's search indexes. Liferay DXP is initially configured to use the embedded HSQL database but the HSQL database is primarily intended for demonstration and trial purposes. Portal property jdbc.default.url sets the Hypersonic embedded HSQL database location.
- deploy: To auto-deploy Liferay DXP plugins, copy them to this folder. Legacy style .war files, 7.0 style .jar files, and .lpkg files from Liferay Marketplace are supported. Portal property auto.deploy.deploy.dir sets the auto-deploy location.
- license: Liferay DXP's copyright and version files are here.
- logs: This folder contains Liferay DXP's log files. The information in Liferay DXP's log files can be quite valuable for system administrators, especially when trying to diagnose a problem. portal-impl.jar's portal-impl/src/META-INF/portal-log4j.xml file sets the location for the log files. To override the log file location, you must use an ext-impl/src/META-INF/portal-log4j-ext.xml file in an Ext plugin.
- osgi: All the JAR files and a few configuration files for Liferay DXP's OSGi runtime belong in this folder. Portal property module.framework.base.dir sets the OSGi folder location. Here are its subfolders:
  * configs: Component configuration files go here
  * core: Liferay DXP's core modules
  * marketplace: Marketplace applications and application suites
  * modules: Modules you've deployed
  * portal: Liferay DXP's non-core modules
  * state: Contains OSGi internal state files for such things as OSGi bundle installation, bundle storage, and more
  * target-platform: Target platform index
  * test: Modules that support test integration
  * war: WAR plugins you've deployed
- patching-tool: (Liferay DXP only) This folder contains patches for Liferay DXP and files for installing the patches.
- tools: For Liferay DXP upgrade and target platform indexer.
- work: Module Jasper work files.

If Liferay DXP is unable to create the resources it needs in the Liferay Home folder or if it finds itself running on certain application servers, it creates a folder called liferay in the home folder of the operating system user that is running Liferay DXP. In this case, the liferay folder becomes Liferay Home. For example, if the operating system user's name is jbloggs, the Liferay Home folder is /home/jbloggs/liferay or C:\Users\jbloggs\liferay.

Extracting a Liferay Bundle

Getting a Liferay DXP bundle up and running involves uncompressing the archive, possibly copying a JDBC driver, and then starting the application server. Let's use the Liferay DXP Tomcat bundle as an example.

1. Extract your Liferay DXP bundle.
2. If you're using a supported open source database or if you're setting up Liferay DXP for demo purposes, you can skip this step. Otherwise, copy your database's JDBC driver .jar file to [Tomcat]/lib/ext.

That's it! You've extracted Liferay DXP, and it's ready for use. This is much easier than doing a manual installation on an app server. If, however, that's what you need to do, please at this point click the link on the left and go through the installation procedure for your app server of choice. When you're finished with the installation (and before you've started Liferay DXP for the first time), come back to this spot, because you must hook it up to your database.

Important: On JDK 11, the setting `-Djava.locale.providers=JRE,COMPAT,CLDR` is required to display four-digit years. Since JDK 9, the Unicode Common Locale Data Repository (CLDR) is the default locales provider. CLDR does not provide years in a four-digit format (see LPS-87191). This setting works around the issue by using JDK 8's default locales provider. Please see Installing Liferay DXP on Tomcat for more information.

Connecting to Your Database

You can connect to your database with JNDI on your app server or the data source shipped with Liferay DXP (recommended). Refer to the article on your app server for instructions on using JNDI. For the internal data source, you can configure it by specifying the configuration in a `portal-ext.properties` file or the Setup Wizard:

1. Create a `portal-ext.properties` file in your Liferay Home folder.
2. Copy a relevant example data source configuration from one of the data source configuration templates or portal properties reference's JDBC section and paste it into the `portal-ext.properties` file.
3. Customize the configuration with the proper host name and user and password credentials for your database, and save the file.

You're ready to start Liferay DXP.

Running for the First Time

Next, start your app server, or start the Liferay DXP app in your app server. Liferay DXP writes log files to folder `[Liferay Home]/logs`.

The first time Liferay DXP starts, it creates all of its database tables. Once it has successfully started, it automatically launches a web browser that displays the Basic Configuration page. If for some reason your browser doesn't load the Basic Configuration page, launch it and navigate to your app server's address and port (for example, http://localhost:8080).

Using the Setup Wizard

The Setup Wizard runs when you start Liferay DXP for the first time. The title of the setup wizard page is Basic Configuration. This page provides a convenient way to make an initial configuration.

There are two sections of the wizard: the portal and the administrator. For the portal, you need to supply the following information:

Portal Name: the name of the portal you're powering with Liferay DXP.

Default Language: choose the default locale of your portal.

Time Zone: select your Liferay DXP instance's default time zone.

For the administrator, you need to supply the following information:

First Name: the first name of the default administrator user
Last Name: the last name of the default administrator user
Email: the email address of the default administrator user

Note: the administrator user's email domain is used as the Liferay DXP instance's default domain (i.e., the `company.default.web.id` portal property).

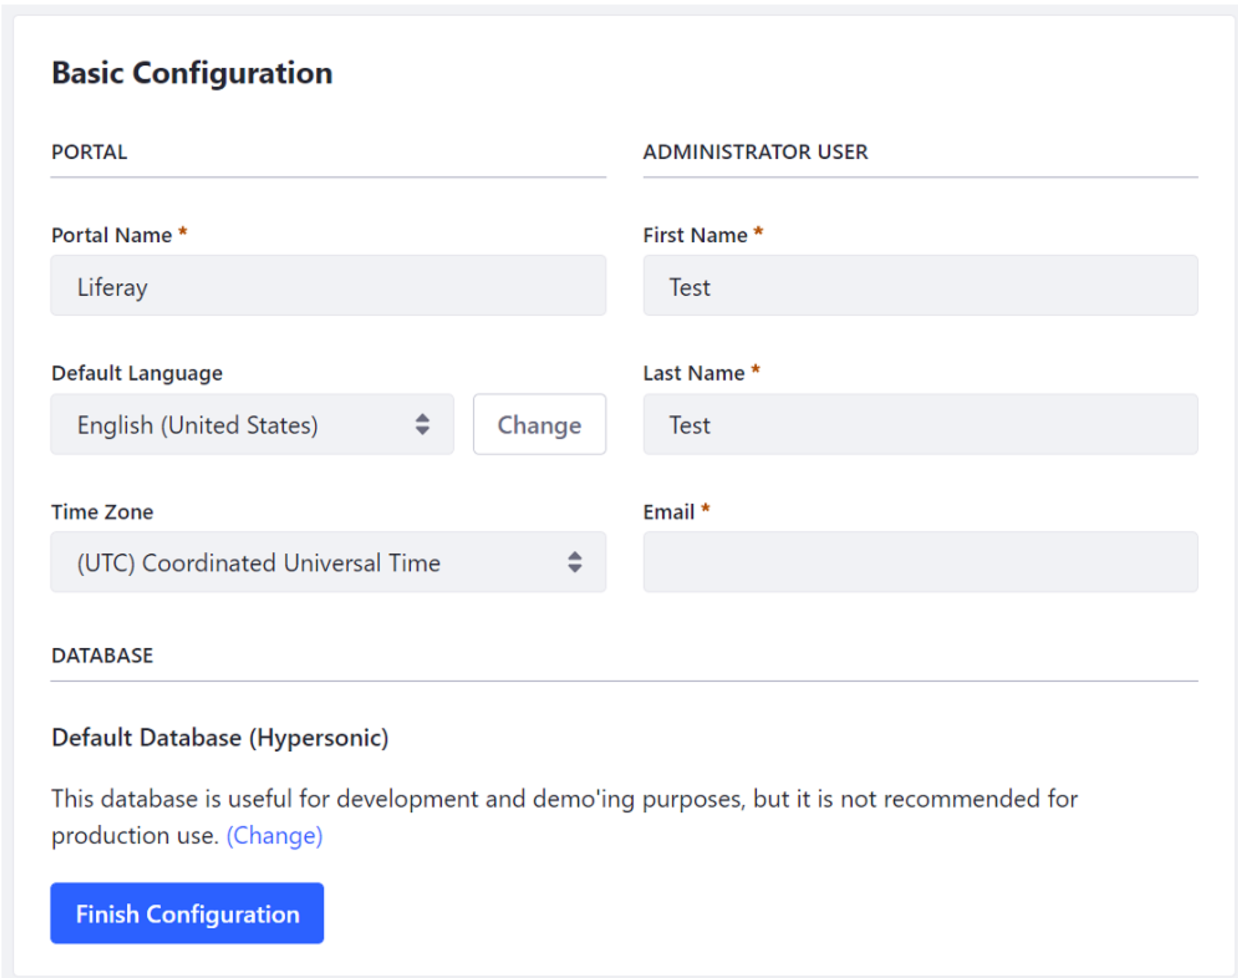

Figure 85.1: Supply the information for your portal and your portal's default administrator user on the Basic Configuration page.

The Basic Configuration page also includes a checkbox labeled *Add Sample Data*. If you check this box, sample data is added to your database. This data includes users, sites, and organizations. The sample data is for demo purposes. If you're installing Liferay DXP on your own machine to explore its features, the sample data may be useful. If, however, you're installing Liferay DXP on a real server, start with a clean system.

Once you've filled out the form, click *Finish Configuration*. The setup wizard creates a `portal-setup-wizard.properties` file which stores the settings that you entered. When you begin customizing your portal's configuration, however, use the `portal-ext.properties` file you created earlier. All the possible properties that can be placed in this file are documented in our reference documentation.

Tip: The wizard is an extremely helpful tool, especially if you're setting up Liferay DXP for the first time. If you're a veteran and you already have your various properties set up, you can disable the setup wizard. If you disable the setup wizard, you must configure everything manually from the `portal-ext.properties` file. To disable the setup wizard, enter `setup.wizard.enabled=false` in your `portal-ext.properties` file. Note that property values in `portal-setup-wizard.properties` (the file created in Liferay Home by the setup wizard) override property values in `portal-ext.properties`.

After you've entered the information requested by the Basic Configuration page, the home page appears. You should set up your mail configuration next.

Configuring Mail

Log in as the administrative user you created in the setup wizard. Click the menu icon and then go to Control Panel → Server Administration → Mail, and have your mail credentials ready.

Fill out the form. You're asked for the following information:

Incoming POP Server: The hostname for a server running the Post Office Protocol. Liferay DXP checks this mailbox for incoming messages, such as message board replies.

Incoming Port: The port on which the POP server is listening.

Use a Secure Network Connection: Use an encrypted connection when connecting to the POP server.

User Name: The user ID Liferay DXP should use to log into the POP server.

Password: The password Liferay DXP should use to log into the POP server.

Outgoing SMTP Server: The hostname for a server running the Simple Mail Transfer Protocol. Liferay DXP uses this server to send emails, such as password change emails and other notifications.

Outgoing Port: The port on which the SMTP server is listening.

Use a Secure Network Connection: Use an encrypted connection when connecting to the SMTP server.

User Name: The user ID Liferay DXP should use to log into the SMTP server.

Password: The password Liferay DXP should use to log into the SMTP server.

Manually specify additional JavaMail properties to override the above configuration: If there are additional properties you need to specify, supply them here.

When you're finished setting up your mail configuration, click *Save*.

Your next step is to convert the search implementation from its default demo mode into a production-ready mode.

85.3 Installing Liferay DXP Manually

The easiest way to install Liferay DXP is to use a bundle. However, this is not always possible. Some organizations have an existing infrastructure into which Liferay DXP must be installed. Other organizations have standardized on a particular application server. Liferay DXP works well with many leading application servers. Before you get started, note that there are two distinct approaches to managing Liferay DXP's data source and mail session. All these topics are covered:

- Using data sources

- Using mail sessions
- Manual configuration for data sources and mail sessions

Start with data sources.

Using Data Sources

Liferay DXP provides two ways to configure your data source:

- Use Liferay DXP's built-in data source
- Use your application's server's JNDI data source

We recommend the built-in data source. Liferay DXP's data source is configured by properties set in a properties file. By default, you can enter database connection information on the Basic Configuration page that appears when Liferay DXP starts for the first time. The Setup Wizard stores the information you entered in a configuration file called `portal-setup-wizard.properties` in your Liferay Home folder. The built-in data source uses this information to connect to the database.

Although using the built-in data source is recommended, that's not the only option. You might prefer to use the data source your application server provides. In this case, a JNDI lookup provides a handle to the data source and the application server manages the connection pools.

To configure your application server's data source, you must create your own configuration file and skip the setup wizard. Since you'd be creating this file *after* the wizard anyway, this isn't a big deal. The Manual Configuration section below demonstrates configuring a JNDI data source.

Since mail sessions are configured similarly to data sources, they're next.

Using Mail Sessions

Liferay DXP uses SMTP to send mail. As with databases, you have two ways to configure your mail server:

- Use Liferay DXP's built-in mail session
- Use your application server's mail session

Using the built-in mail session is recommended. After you've started Liferay DXP, you can configure a mail server through the Control Panel. The default configuration is a mail server on the same machine running Liferay DXP. If this is not your configuration, you must modify the defaults. To do this, use a `portal-ext.properties` file in your Liferay Home folder (see below).

To use your application server's mail session, you must create it in your application server. Once you've created a mail session, you can point Liferay DXP to it through your `portal-ext.properties` file or through the Control Panel.

If you plan to use Liferay DXP to manage both your database connection and mail session, enter your database connection information on the Basic Configuration page when Liferay DXP first starts, and then enter your mail server information through the Control Panel.

If you plan to let your application server manage your database connection or your mail server, you can't use Liferay DXP's setup wizard: you must follow the instructions in the Manual Configuration section below.

The installation articles for each application server also include instructions for configuring your application server to manage the database connection and mail server.

Manual Configuration

To have your application server manage your database connection or mail server (or both), you must manually create this configuration. Create a text file called `portal-ext.properties` in your Liferay Home folder. This file overrides the default properties.

To use your application server's data source, create a connection pool in your application server that points to your database. The connection pool should be called `jdbc/LiferayPool`. This is spelled out for each application server in its article. To tell Liferay DXP to use your `jdbc/LiferayPool` connection pool, add the following directive to your `portal-ext.properties` file:

```
jdbc.default.jndi.name=jdbc/LiferayPool
```

Next, install Liferay DXP according to the article for your application server. Once it's installed, you can set up the mail configuration.

You should use the Control Panel to create the mail configuration. Go to *Control Panel → Configuration → Server Administration → Mail* and enter your settings for your mail session settings.

You can also configure this with the `portal-ext.properties` file, which lets you do the configuration once and then copy the configuration file to multiple machines. To use the built-in mail session, use the following properties and customize their values for your environment:

```
mail.session.mail.pop3.host=localhost
mail.session.mail.pop3.password=
mail.session.mail.pop3.port=110
mail.session.mail.pop3.user=
mail.session.mail.smtp.auth=false
mail.session.mail.smtp.host=localhost
mail.session.mail.smtp.password=
mail.session.mail.smtp.port=25
mail.session.mail.smtp.user=
mail.session.mail.store.protocol=pop3
mail.session.mail.transport.protocol=smtp
```

To use your application server's mail session, create it first. Then specify it in the `portal-ext.properties` file:

```
mail.session.jndi.name=mail/MailSession
```

When you're finished, save the file.

Logging

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass gatherExternProperties
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass checkControlFlow
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patter
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
```

85.4 Installing Liferay DXP on Tomcat

7.0 bundled with Tomcat 9 is available on the Help Center (DXP) or Liferay Downloads (Portal CE). The Tomcat bundle contains JARs, scripts, and configuration files required for installing Liferay DXP on a clean Tomcat 9 application server. Copying these files from a Liferay DXP Tomcat bundle facilitates installing Liferay DXP on Tomcat.

Whether you copy bundle files (recommended) or download and create the files, you must download these *Additional Files* for DXP or Portal CE:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Liferay DXP requires a Java JDK 8 or 11.

Note: Please see the compatibility matrix for information on supported JDKs, databases, and environments.

Here are the basic steps for installing Liferay DXP on Tomcat:

- Installing dependencies to your application server
- Configuring your application server for Liferay DXP
- Deploying the Liferay DXP WAR file to your application server

Liferay Home is the folder containing your Tomcat server folder. After installing and deploying Liferay DXP, Liferay Home contains `data`, `deploy`, `license`, and `osgi` folders. `$TOMCAT_HOME` refers to your Tomcat server folder. It is usually named `tomcat-[version]` or `apache-tomcat-[version]`.

Installing Dependencies

Liferay DXP depends on many JARs included by Liferay DXP Tomcat bundle. Some of the bundle's JARs are not strictly required but can still be useful. If you don't have a bundle, download the required JARs from third-parties, as described below.

1. Unzip the Dependencies ZIP file contents in the `$TOMCAT_HOME/lib/ext` folder (create this folder if it doesn't exist).

2. Download a database driver `.jar` file and copy it to the `$CATALINA_BASE/lib/ext` folder. For a list of supported databases, see Liferay's compatibility matrix

3. Download the following JARs or copy them from a Liferay DXP Tomcat bundle to the `$TOMCAT_HOME/lib/ext` folder:

    - `activation.jar`
    - `ccpp.jar`
    - `jms.jar`
    - `jta.jar`
    - `jutf7.jar`
    - `mail.jar`

    - `persistence.jar`
    - `support-tomcat.jar`

4. Unzip the OSGi Dependencies ZIP file contents in the `[Liferay Home]/osgi` folder (create this folder if it doesn't exist).

Configuring Tomcat

Configuring Tomcat to run Liferay DXP includes these things:

- Setting environment variables
- Specifying a web application context for Liferay DXP
- Setting properties and descriptors

Optionally, you can configure Tomcat to manage these things for Liferay DXP:

- Data source
- Mail session

Start with configuring Tomcat to run Liferay DXP.

1. If you have a Liferay DXP Tomcat bundle, copy the `setenv.bat` and `setenv.sh` files from it to your `$CATALINA_BASE/bin` folder. If not, create these scripts.

   The scripts set JVM options for Catalina, which is Tomcat's servlet container. Among these options is the location of the Java runtime environment. If this environment is not available on your server globally, you must set its location in in these files so Tomcat can run. Do this by pointing the `JAVA_HOME` environment variable to a Liferay DXP-supported JRE:

   ```
   export JAVA_HOME=/usr/lib/jvm/java-8-jdk
   export PATH=$JAVA_HOME/bin:$PATH
   ```

   Then configure Catalina's JVM options to support Liferay DXP.

   Unix:

   ```
   CATALINA_OPTS="$CATALINA_OPTS -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Dorg.apache.catalina.loader.WebappClassLoader.ENABLE_CLE
   Duser.timezone=GMT -Xmx2048m -XX:MaxMetaspaceSize=512m"
   ```

   Windows:

   ```
   set "CATALINA_OPTS=%CATALINA_OPTS% -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Dorg.apache.catalina.loader.WebappClassLoader.ENABL
   Duser.timezone=GMT -Xmx2048m -XX:MaxMetaspaceSize=512m"
   ```

   This sets the file encoding to UTF-8, prefers an IPv4 stack over IPv6, prevents Tomcat from working around garbage collection bugs relating to static or final fields (these bugs don't exist in Liferay DXP and working around them causes problems with the logging system), sets the time zone to GMT, gives the JVM 2GB of RAM, and limits Metaspace to 500MB.

```
 **Important:** For Liferay DXP to work properly, the application server JVM
 must use the `GMT` time zone and `UTF-8` file encoding.
```

```
 **Important:** On JDK 11, the setting `-Djava.locale.providers=JRE,COMPAT,CLDR` is required to display four-digit years. Since JDK 9, the Unicode Co
digit format (see [LPS-87191](https://issues.liferay.com/browse/LPS-87191)). This setting works around the issue by using JDK 8's default locales pro
```

```
After installation, tune your system (including these JVM options) for
performance.
```

2. If you have a Liferay DXP Tomcat bundle, copy its `$CATALINA_BASE/conf/Catalina/localhost/ROOT.xml` file to the corresponding location in your application server. Create the file path if it doesn't exist. If you don't have a Liferay DXP Tomcat bundle, create a `ROOT.xml` file.

   The `ROOT.xml` file specifies a web application context for Liferay DXP. `ROOT.xml` looks like this:

```
<Context crossContext="true" path="">

    <!-- JAAS -->

    <!--<Realm
        className="org.apache.catalina.realm.JAASRealm"
        appName="PortalRealm"
        userClassNames="com.liferay.portal.kernel.security.jaas.PortalPrincipal"
        roleClassNames="com.liferay.portal.kernel.security.jaas.PortalRole"
    />-->

    <!--
    Uncomment the following to disable persistent sessions across reboots.
    -->

    <!--<Manager pathname="" />-->

    <!--
    Uncomment the following to not use sessions. See the property
    "session.disabled" in portal.properties.
    -->

    <!--<Manager className="com.liferay.support.tomcat.session.SessionLessManagerBase" />-->

    <Resources>
        <PreResources
            base="${catalina.base}/lib/ext/portal"
            className="com.liferay.support.tomcat.webresources.ExtResourceSet"
            webAppMount="/WEB-INF/lib"
        />
    </Resources>
</Context>
```

   Setting `crossContext="true"` lets multiple web applications use the same class loader. This configuration includes commented instructions and tags for configuring a JAAS realm, disabling persistent sessions, and disabling sessions entirely.

3. Provide Catalina access to the JARs in `$CATALINA_BASE/lib/ext` by opening your `$CATALINA_BASE/conf/catalina.p` file and appending this value to the `common.loader` property:

```
,"${catalina.home}/lib/ext/global","${catalina.home}/lib/ext/global/*.jar","${catalina.home}/lib/ext","${catalina.home}/lib/ext/*.jar"
```

4. Make sure to use UTF-8 URI encoding consistently. If you have a Liferay DXP Tomcat bundle, copy the `$CATALINA_BASE/conf/server.xml` file to your server. If not, open your `$CATALINA_BASE/conf/server.xml` file and add the attribute `URIEncoding="UTF-8"` to HTTP and AJP connectors that use `redirectPort=8443`. Here are examples:

   Old:

   ```
   <Connector port="8080" protocol="HTTP/1.1" connectionTimeout="20000" redirectPort="8443" />
   ```

   New:

   ```
   <Connector port="8080" protocol="HTTP/1.1" connectionTimeout="20000" redirectPort="8443" URIEncoding="UTF-8" />
   ```

   Old:

   ```
   <Connector port="8009" protocol="AJP/1.3" redirectPort="8443" />
   ```

   New:

   ```
   <Connector port="8009" protocol="AJP/1.3" redirectPort="8443" URIEncoding="UTF-8" />
   ```

5. If you're on Unix, Linux, or Mac OS, make the shell scripts in your `$CATALINA_HOME/bin` and `$CATALINA_BASE/bin` folders executable by running this command in each folder:

   `chmod a+x *.sh`

Checkpoint:
Your application server is configured to run Liferay DXP.

1. The file encoding, user time-zone, and preferred protocol stack are set in `setenv.sh`.

2. The default memory available and Metaspace limit are set.

3. `$CATALINA_BASE/conf/Catalina/localhost/ROOT.xml` declares the web application context.

4. The `common.loader` property in `$CATALINA_BASE/conf/catalina.properties` grants Catalina access to the JARs in `$CATALINA_BASE/lib/ext`.

5. `$CATALINA_BASE/conf/server.xml` sets UTF-8 encoding.

6. The scripts in Tomcat's `bin` folders are executable.

Database Configuration

The easiest way to handle your database configuration is to let Liferay DXP manage your data source. Liferay DXP's Basic Configuration page lets you configure Liferay DXP's built-in data source. If you want to use the built-in data source, skip this section.

If you want Tomcat to manage your data source, follow these steps:

1. Make sure your database server is installed and working. If it's installed on a different machine, make sure it's accessible from your Liferay DXP machine.

2. Open `$CATALINA_BASE/conf/Catalina/localhost/ROOT.xml` and add your data source as a `Resource` in your web application `Context`:

```
<Context...>
    ...
    <Resource
        name="jdbc/LiferayPool"
        auth="Container"
        type="javax.sql.DataSource"
        driverClassName="com.mysql.jdbc.Driver"
        url="jdbc:mysql://localhost/lportal?useUnicode=true&amp;characterEncoding=UTF-8"
        username="root"
        password="root"
        maxTotal="100"
        maxIdle="30"
        maxWaitMillis="10000"
    />
</Context>
```

   The resource definition above is for a MySQL database named `lportal` that has a user named `root` whose password is `root`. Replace these values with your own.

3. In a `portal-ext.properties` file in your Liferay Home, specify your data source:

```
jdbc.default.jndi.name=jdbc/LiferayPool
```

You created a data source for Tomcat to manage and configured Liferay DXP to use it. Mail session configuration is next.

Mail Configuration

As with database configuration, the easiest way to configure mail is to let Liferay DXP handle your mail session. If you want to use Liferay DXP's built-in mail session, skip this section and configure the mail session in the Control Panel.

If you want to manage your mail session with Tomcat, follow these steps:

1. Open `$CATALINA_BASE/conf/Catalina/localhost/ROOT.xml` and add your mail session as a `Resource` in your web application `Context`. Make sure to replace the example mail session values with your own.

```
<Context...>
    ...
    <Resource
        name="mail/MailSession"
        auth="Container"
```

```
        type="javax.mail.Session"
        mail.pop3.host="pop.gmail.com"
        mail.pop3.port="110"
        mail.smtp.host="smtp.gmail.com"
        mail.smtp.port="465"
        mail.smtp.user="user"
        mail.smtp.password="password"
        mail.smtp.auth="true"
        mail.smtp.starttls.enable="true"
        mail.smtp.socketFactory.class="javax.net.ssl.SSLSocketFactory"
        mail.imap.host="imap.gmail.com"
        mail.imap.port="993"
        mail.transport.protocol="smtp"
        mail.store.protocol="imap"
    />
</Context>
```

2. In your `portal-ext.properties` file in Liferay Home, reference your mail session:

```
mail.session.jndi.name=mail/MailSession
```

You've created a mail session for Tomcat to manage and configured Liferay DXP to use it.

Deploying Liferay

Now you're ready to deploy Liferay DXP using the Liferay DXP WAR file.

1. If you are manually installing Liferay DXP on a clean Tomcat server, delete the contents of the `$CATALINA_BASE/webapps/ROOT` folder. This removes the default Tomcat home page.

2. Extract the Liferay DXP `.war` file to `$CATALINA_BASE/webapps/ROOT`.

   It's time to launch Liferay DXP on Tomcat!

3. Start Tomcat by navigating to `$CATALINA_HOME/bin` and executing `./startup.sh`. Alternatively, execute `./catalina.sh run` to tail Liferay DXP's log file. The log audits startup activities and is useful for debugging deployment.

Congratulations on successfully installing and deploying Liferay DXP on Tomcat!

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass gatherExternProperties
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass checkControlFlow
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patter
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
 current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patte
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
```

85.5 Installing Liferay DXP on Wildfly

7.0 bundled with Wildfly 11 is available on the Help Center (DXP) or Liferay Downloads (Portal CE). 7.0 supports deployment to Wildfly 10 and Wildfly 11. Even if you want to manually install Liferay DXP on an existing Wildfly application server, it can be helpful to download a Liferay DXP Wildfly bundle to make gathering the dependencies easier. Before proceeding, also download these *Additional Files* for DXP or Portal CE:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Liferay DXP requires a Java JDK 8 or 11.

Note: Please see the compatibility matrix for information on supported JDKs, databases, and environments.

Installing Liferay DXP manually takes three steps:

- Installing dependencies to your application server
- Configuring your application server for Liferay DXP
- Deploying the Liferay DXP WAR file to your application server

Liferay Home is the folder containing your Wildfly server folder. After installing and deploying Liferay DXP, the Liferay Home folder contains the Wildfly server folder as well as `data`, `deploy`, `logs`, and `osgi` folders. `$WILDFLY_HOME` refers to your Wildfly server folder. It is usually named `wildfly-[version]`.

Installing Dependencies

Liferay DXP depends on many JARs that are included in the Liferay DXP Wildfly bundle. Some of the bundle's JARs are not strictly required but can still be useful. If you don't have a Liferay DXP Wildfly bundle, download the required JARs from third-parties as described below.

1. Create the folder `$WILDFLY_HOME/modules/com/liferay/portal/main` if it doesn't exist and extract the dependencies ZIP JARs to it.

2. Download your database driver `.jar` file and copy it into the same folder. Please see the compatibility matrix for a list of supported databases.

3. Create the file `module.xml` in the `$WILDFLY_HOME/modules/com/liferay/portal/main` folder. In the file, declare the portal module and all of its required resources and dependencies:

```
<?xml version="1.0"?>

<module xmlns="urn:jboss:module:1.0" name="com.liferay.portal">
    <resources>
        <resource-root path="[place your database vendor's JAR file name here]" />
        <resource-root path="[place a Liferay dependencies JAR file name here]" />
        <!-- Add a resource-root element for each Liferay dependencies JAR -->
    </resources>
```

```
    <dependencies>
        <module name="javax.api" />
        <module name="javax.mail.api" />
        <module name="javax.servlet.api" />
        <module name="javax.servlet.jsp.api" />
        <module name="javax.transaction.api" />
    </dependencies>
</module>
```

Replace `[place your database vendor's JAR file name here]` with the driver JAR for your database.

For each JAR in the Liferay dependencies ZIP, add a `resource-root` element with its `path` attribute set to the JAR name. For example, add a `resource-root` element like this for the `com.liferay.petra.concurrent.jar` file:

```
<resource-root path="com.liferay.petra.concurrent.jar" />
```

4. Create an `osgi` folder in your Liferay Home folder. Extract the OSGi ZIP file that you downloaded into the `osgi` folder.

   The `osgi` folder provides the necessary modules for Liferay DXP's OSGi runtime.

Checkpoint:

1. The contents of the Dependencies zip have been placed in the `$WILDFLY_HOME/modules/com/liferay/portal/main` folder:
2. Your database vendor's JDBC driver has been placed in `$WILDFLY_HOME/modules/com/liferay/portal/main` folder and listed as a dependency.
3. The `module.xml` has listed all JARs in the `<resource-root>` elements.
4. The OSGi dependencies have been unzipped in the `osgi` folder located inside the `${Liferay.home}` folder.

Running Liferay DXP on Wildfly in Standalone Mode vs. Domain Mode

Wildfly can be launched in either *standalone* mode or *domain* mode. Domain mode allows multiple application server instances to be managed from a single control point. A collection of such application servers is known as a *domain*. For more information on standalone mode vs. domain mode, please refer to the section on this topic in the Wildfly Admin Guide. Liferay DXP fully supports Wildfly in standalone mode but not in domain mode.

You can run Liferay DXP on Wildfly in domain mode, but this method is not fully supported. In particular, Liferay DXP's hot-deploy does not work with a managed deployment, since Wildfly manages the content of a managed deployment by copying files (exploded or non-exploded). This prevents JSP hooks and Ext plugins from working as intended. For example, JSP hooks don't work on Wildfly running in managed domain mode, since Liferay DXP's JSP override mechanism relies on the application server. Since both of these features are deprecated, however, you may not be using them.

The command line interface is recommended for domain mode deployments.

Note: This does not prevent Liferay DXP from running in a clustered environment on multiple Wildfly servers. You can set up a cluster of Liferay DXP instances running on Wildfly servers running in standalone mode. Please refer to the chapter of this guide on Liferay DXP Clustering for information on setting up a Liferay DXP cluster.

Configuring Wildfly

Configuring Wildfly to run Liferay DXP includes these things:

- Setting environment variables
- Setting properties and descriptors
- Removing unnecessary configurations

Optionally, you can configure Wildfly to manage Liferay DXP's data source and mail session.
Start with configuring Wildfly to run Liferay DXP.
Make the following modifications to `$WILDFLY_HOME/standalone/configuration/standalone.xml`:

1. In the `<jsp-config>` tag, set the Java VM compatibility for Liferay source and class files. They are compatible with Java 8 by default.

```
<jsp-config development="true" source-vm="1.8" target-vm="1.8" />
```

2. Locate the closing `</extensions>` tag. Directly beneath that tag, insert the following system properties:

```
<system-properties>
    <property name="org.apache.catalina.connector.URI_ENCODING" value="UTF-8" />
    <property name="org.apache.catalina.connector.USE_BODY_ENCODING_FOR_QUERY_STRING" value="true" />
</system-properties>
```

3. Add the following `<filter-spec>` tag within the `<console-handler>` tag, directly below the `<level name="INFO"/>` tag:

```
<filter-spec value="not(any(match(&quot;WFLYSRV0059&quot;),match(&quot;WFLYEE0007&quot;)))" />
```

4. Add a timeout for the deployment scanner by setting `deployment-timeout="360"` as seen in the excerpt below.

```
<subsystem xmlns="urn:jboss:domain:deployment-scanner:2.0">
    <deployment-scanner deployment-timeout="360" path="deployments" relative-to="jboss.server.base.dir" scan-interval="5000" runtime-
failure-causes-rollback="${jboss.deployment.scanner.rollback.on.failure:false}"/>
</subsystem>
```

5. Add the following JAAS security domain to the security subsystem `<security-domains>` defined in element `<subsystem xmlns="urn:jboss:domain:security:2.0">`.

```
<security-domain name="PortalRealm">
    <authentication>
        <login-module code="com.liferay.portal.kernel.security.jaas.PortalLoginModule" flag="required" />
    </authentication>
</security-domain>
```

6. Remove the two code snippets providing welcome content:

```
<location name="/" handler="welcome-content"/>
```

and

```
<handlers>
    <file name="welcome-content" path="${jboss.home.dir}/welcome-content"/>
</handlers>
```

Checkpoint:
Before continuing, verify the following properties have been set in the `standalone.xml` file:

1. The new `<system-property>` is added.
2. The new `<filter-spec>` is added.
3. The `<deployment-timeout>` is set to `360`.
4. The new `<security-domain>` is created.
5. Welcome content is removed.
6. The `<jsp-config>` tag contains its new attributes.

Now you must configure your JVM and startup scripts.

In the `$WILDFLY_HOME/bin/` folder, you must make these modifications to your standalone domain's configuration script file `standalone.conf` (`standalone.conf.bat` on Windows):

- Set the file encoding
- Set the user time-zone
- Set the preferred protocol stack
- Increase the default amount of memory available.

Make the following edits as applicable for your operating system:
Windows:

1. Comment out the initial `JAVA_OPTS` assignment like this:

```
rem set "JAVA_OPTS=-Xms64M -Xmx512M -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=256m"
```

2. Add the following `JAVA_OPTS` assignment one line above the `:JAVA_OPTS_SET` line found at end of the file:

```
set "JAVA_OPTS=%JAVA_OPTS% -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Djboss.as.management.blocking.timeout=480 -
Duser.timezone=GMT -Xmx2048m -XX:MaxMetaspaceSize=512m -XX:MetaspaceSize=200m"
```

Unix:

1. Below the `if [ "x$JAVA_OPTS" = "x" ];` statement, replace this `JAVA_OPTS` statement:

```
JAVA_OPTS="-Xms64m -Xmx512m -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=256m -Djava.net.preferIPv4Stack=true"
```

with this:

```
JAVA_OPTS="-Djava.net.preferIPv4Stack=true"
```

2. Add the following statement to the bottom of the file:

```
JAVA_OPTS="$JAVA_OPTS -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true  -Djboss.as.management.blocking.timeout=480 -
Duser.timezone=GMT -Xmx2048m -XX:MaxMetaspaceSize=512m -XX:MetaspaceSize=200m"
```

Important: For Liferay DXP to work properly, the application server JVM must use the `GMT` time zone and `UTF-8` file encoding.

Important: On JDK 11, the setting `-Djava.locale.providers=JRE,COMPAT,CLDR` is required to display four-digit years. Since JDK 9, the Unicode Common Locale Data Repository (CLDR) is the default locales provider. CLDR does not provide years in a four-digit format (see LPS-87191). This setting works around the issue by using JDK 8's default locales provider.

Note: If you plan on using the IBM JDK with your Wildfly server, you must complete some additional steps. First, navigate to the `$WILDFLY_HOME/modules/com/liferay/portal/main/module.xml` file and insert the following dependency within the `<dependencies>` element:

```
<module name="ibm.jdk" />
```

Then navigate to the `$WILDFLY_HOME/modules/system/layers/base/sun/jdk/main/module.xml` file and insert the following path names inside the `<paths>...</paths>` element:

```
<path name="com/sun/crypto" />
<path name="com/sun/crypto/provider" />
<path name="com/sun/org/apache/xml/internal/resolver" />
<path name="com/sun/org/apache/xml/internal/resolver/tools" />
```

The added paths resolve issues with deployment exceptions and image uploading problems.

Checkpoint:
At this point, you've finished configuring the application server's JVM settings.

1. The file encoding, user time-zone, preferred protocol stack have been set in the `JAVA_OPTS` in the `standalone.conf.bat` file.

2. The default amount of memory available has been increased.

The prescribed script modifications are now complete for your Liferay DXP installation on Wildfly. Next you'll configure your database.

Database Configuration

The easiest way to handle database configuration is to let Liferay DXP manage your data source. The Basic Configuration page lets you configure Liferay DXP's built-in data source. If you want to use the built-in data source, skip this section.

If you want Wildfly to manage your data source, follow these steps:

1. Add the data source inside the `$WILDFLY_HOME/standalone/configuration/standalone.xml` file's the `<datasources>` element.

```
<datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-java-
context="true" use-ccm="true">
    <connection-url>[place the URL to your database here]</connection-url>
    <driver>[place the driver name here]</driver>
    <security>
        <user-name>[place your user name here]</user-name>
        <password>[place your password here]</password>
    </security>
</datasource>
```

Make sure to replace the database URL, user name, and password with the appropriate values.

```
**Note:** If you must change your datasource `jndi-name` to something
different, you must also edit the `datasource` element in the
`<default-bindings>` tag.
```

---

2. Add the driver to the `standalone.xml` file's `<drivers>` element also found within the `<datasources>` element.

```
<drivers>
    <driver name="[name of driver must match name above]" module="com.liferay.portal">
        <driver-class>[place your JDBC driver class here]</driver-class>
    </driver>
</drivers>
```

A final data source subsystem that uses MySQL should look like this:

```
<subsystem xmlns="urn:jboss:domain:datasources:5.0">
    <datasources>
        <datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-
java-context="true" use-ccm="true">
            <connection-url>jdbc:mysql://localhost/lportal</connection-url>
            <driver>mysql</driver>
            <security>
                <user-name>root</user-name>
                <password>root</password>
            </security>
        </datasource>
        <drivers>
            <driver name="mysql" module="com.liferay.portal"/>
        </drivers>
    </datasources>
</subsystem>
```

3. In a portal-ext.properties file in your Liferay Home, specify your data source:

```
jdbc.default.jndi.name=java:jboss/datasources/ExampleDS
```

Now that you've configured your data source, the mail session is next.

Mail Configuration

As with database configuration, the easiest way to configure mail is to let Liferay DXP handle your mail session. If you want to use Liferay DXP's built-in mail session, skip this section and configure the mail session in the Control Panel.

If you want to manage your mail session with Wildfly, follow these steps:

1. Specify your mail subsystem in the $WILDFLY_HOME/standalone/configuration/standalone.xml file like this:

```
<subsystem xmlns="urn:jboss:domain:mail:3.0">
    <mail-session jndi-name="java:jboss/mail/MailSession" name="mail-smtp">
        <smtp-server ssl="true" outbound-socket-binding-ref="mail-smtp" username="USERNAME" password="PASSWORD"/>
   </mail-session>
</subsystem>
...
<socket-binding-group name="standard-sockets" default-interface="public" port-offset="${jboss.socket.binding.port-
offset:0}">
...
<outbound-socket-binding name="mail-smtp">
        <remote-destination host="smtp.gmail.com" port="465"/>
    </outbound-socket-binding>
</socket-binding-group>
```

2. In your portal-ext.properties file in Liferay Home, reference your mail session:

```
mail.session.jndi.name=java:jboss/mail/MailSession
```

You've got mail! Next, you'll deploy Liferay DXP to your Wildfly app server.

Deploying Liferay DXP

Now you're ready to deploy Liferay DXP using the Liferay DXP WAR file.

1. If the folder $WILDFLY_HOME/standalone/deployments/ROOT.war already exists in your Wildfly installation, delete all of its subfolders and files. Otherwise, create a new folder called $WILDFLY_HOME/standalone/deployments/ROOT.war.

2. Unzip the Liferay DXP .war file into the ROOT.war folder.

3. To trigger deployment of ROOT.war, create an empty file named ROOT.war.dodeploy in your $WILDFLY_HOME/standalone/deployments/ folder. On startup, Wildfly detects the presence of this file and deploys it as a web application.

4. Start the Wildfly application server by navigating to $WILDFLY_HOME/bin and running standalone.bat or standalone.sh.

Congratulations; you've deployed Liferay DXP on Wildfly!

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass gatherExternProperties
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass checkControlFlow
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patter
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patte
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
```

85.6 Installing Liferay DXP on JBoss EAP

Installing Liferay DXP on JBoss EAP 7.1 takes three steps:

- Installing Liferay DXP dependencies to your application server
- Configuring your application server for Liferay DXP
- Installing the Liferay DXP WAR file to your application server

Before proceeding, download these files for DXP or Portal CE:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi JARs ZIP file

Liferay DXP requires a Java JDK 8 or 11.

Note: Please see the compatibility matrix for information on supported JDKs, databases, and environments.

Liferay Home is the folder containing your JBoss server folder. After installing and deploying Liferay DXP, the Liferay Home folder contains the JBoss server folder as well as `data`, `deploy`, `logs`, and `osgi` folders. `$JBOSS_HOME` refers to your JBoss server folder. This folder is usually named `jboss-eap-[version]`.

Installing Dependencies

Liferay DXP depends on several Liferay-specific and third-party JARs. Download and install the required JARs as described below.

1. Create the folder `$JBOSS_HOME/modules/com/liferay/portal/main` if it doesn't exist and extract the Liferay dependencies ZIP JARs to it.

2. Download your database driver `.jar` file and copy it into the same folder. Please see the compatibility matrix for a list of supported databases.

3. Create the file `module.xml` in the `$JBOSS_HOME/modules/com/liferay/portal/main` folder. In the file, declare the portal module and all of its required resources and dependencies:

```
<?xml version="1.0"?>

<module xmlns="urn:jboss:module:1.0" name="com.liferay.portal">
    <resources>
        <resource-root path="[place your database vendor's JAR file name here]" />
        <resource-root path="[place a Liferay dependencies JAR file name here]" />
        <!-- Add a resource-root element for each Liferay dependencies JAR -->
    </resources>
    <dependencies>
        <module name="javax.api" />
        <module name="javax.mail.api" />
        <module name="javax.servlet.api" />
        <module name="javax.servlet.jsp.api" />
        <module name="javax.transaction.api" />
    </dependencies>
</module>
```

   Replace `[place your database vendor's JAR file name here]` with the driver JAR for your database.

   For each JAR in the Liferay dependencies ZIP, add a `resource-root` element with its `path` attribute set to the JAR name. For example, add a `resource-root` element like this for the `com.liferay.petra.concurrent.jar` file:

```
<resource-root path="com.liferay.petra.concurrent.jar" />
```

4. Create an `osgi` folder in your Liferay Home folder. Extract the OSGi ZIP file that you downloaded into the `osgi` folder.

   The `osgi` folder provides the necessary modules for Liferay DXP's OSGi runtime.

Checkpoint:

1. The dependencies files have been unzipped into the `$JBOSS_HOME/modules/com/liferay/portal/main` folder and a database jar.
2. The `module.xml` contains all JARs in the `<resource-root-path>` elements.
3. The `osgi` dependencies have been unzipped into the `osgi` folder.

Running Liferay DXP on JBoss EAP in Standalone Mode vs. Domain Mode

JBoss EAP can be launched in either *standalone* mode or *domain* mode. Domain mode allows multiple application server instances to be managed from a single control point. A collection of such application servers is known as a *domain*. For more information on standalone mode vs. domain mode, please refer to the section on this topic in the JBoss EAP Product Documentation.

Liferay DXP supports JBoss EAP when it runs in standalone mode but not when it runs in domain mode. Liferay DXP's hot-deploy does not work with a managed deployment, since JBoss manages the content of a managed deployment by copying files (exploded or non-exploded). This prevents JSP hooks and Ext plugins from working as intended. For example, JSP hooks don't work on JBoss EAP running in managed domain mode, since Liferay DXP's JSP override mechanism relies on the application server. Since both of these features are deprecated, however, you may not be using them.

The command line interface is recommended for domain mode deployments.

Note: This does not prevent Liferay DXP from running in a clustered environment on multiple JBoss servers. You can set up a cluster of Liferay DXP instances running on JBoss EAP servers running in standalone mode. Please refer to the Liferay DXP Clustering section for information on setting up a Liferay DXP cluster.

Configuring JBoss

Configuring JBoss to run Liferay DXP includes these things:

- Setting environment variables
- Setting properties and descriptors
- Removing unnecessary configurations

Optionally, you can configure JBoss to manage these things for Liferay DXP:

- Data source
- Mail session

Start with configuring JBoss to run Liferay DXP.
Make the following modifications to `$JBOSS_HOME/standalone/configuration/standalone.xml`:

1. In the `<jsp-configuration>` tag, set the Java VM compatibility for Liferay source and class files. They are compatible with Java 8 by default.

```
<jsp-configuration development="true" source-vm="1.8" target-vm="1.8" />
```

2. Locate the closing `</extensions>` tag. Directly beneath that tag, insert these system properties:

```
<system-properties>
    <property name="org.apache.catalina.connector.URI_ENCODING" value="UTF-8" />
    <property name="org.apache.catalina.connector.USE_BODY_ENCODING_FOR_QUERY_STRING" value="true" />
</system-properties>
```

3. Add the following `<filter-spec>` tag within the `<console-handler>` tag, directly below the `<level name="INFO"/>` tag:

```
<filter-spec value="not(any(match(&quot;WFLYSRV0059&quot;),match(&quot;WFLYEE0007&quot;)))" />
```

4. Add a timeout for the deployment scanner by setting `deployment-timeout="360"` as seen in the excerpt below.

```
<subsystem xmlns="urn:jboss:domain:deployment-scanner:2.0">
    <deployment-scanner deployment-timeout="360" path="deployments" relative-to="jboss.server.base.dir" scan-interval="5000"/>
</subsystem>
```

5. Add the following JAAS security domain to the security subsystem `<security-domains>` defined in element `<subsystem xmlns="urn:jboss:domain:security:2.0">`.

```
<security-domain name="PortalRealm">
    <authentication>
        <login-module code="com.liferay.portal.kernel.security.jaas.PortalLoginModule" flag="required" />
    </authentication>
</security-domain>
```

6. Remove the two code snippets providing welcome content:

   ```
   <location name="/" handler="welcome-content"/>
   ```

   and

   ```
   <handlers>
       <file name="welcome-content" path="${jboss.home.dir}/welcome-content"/>
   </handlers>
   ```

7. Find the `<jsp-config/>` tag and set the `development`, `source-vm`, and `target-vm` attributes in the tag. Once finished, the tag should look like this:

   ```
   <jsp-config development="true" source-vm="1.8" target-vm="1.8"/>
   ```

Checkpoint:
Before continuing, verify the following properties have been set in the `standalone.xml` file:

1. The new `<system-property>` is added.
2. The new `<filter-spec>` is added.
3. The `<deployment-timeout>` is set to `360`.
4. The new `<security-domain>` is created.
5. Welcome content is removed.
6. The `<jsp-config>` tag contains its new attributes.

Now you should configure your JVM and startup scripts.

In the `$JBOSS_HOME/bin` folder, you must make these modifications to your standalone domain's configuration script file `standalone.conf` (`standalone.conf.bat` on Windows):

- Set the file encoding to `UTF-8`
- Set the user time zone to `GMT`
- Set the preferred protocol stack
- Increase the default amount of memory available.

Important: For Liferay DXP to work properly, the application server JVM must use the `GMT` time zone and `UTF-8` file encoding.

Make the following edits as applicable to your operating system:

Windows

1. Comment out the initial `JAVA_OPTS` assignment as demonstrated in the following line:

   ```
   rem set "JAVA_OPTS=-Xms1G -Xmx1G -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=256m"
   ```

2. Add the following `JAVA_OPTS` assignment one line above the `:JAVA_OPTS_SET` line found at end of the file:

   ```
   set "JAVA_OPTS=%JAVA_OPTS% -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Djboss.as.management.blocking.timeout=480 -
   Duser.timezone=GMT -Xmx2048m -XX:MaxMetaspaceSize=384m"
   ```

Unix

1. Below the `if [ "x$JAVA_OPTS" = "x" ];` statement, replace this `JAVA_OPTS` statement:

   ```
   JAVA_OPTS="-Xms1303m -Xmx1303m -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=256m -Djava.net.preferIPv4Stack=true
   ```

   with this:

   ```
   JAVA_OPTS="-Djava.net.preferIPv4Stack=true"
   ```

2. Add the following statement to the bottom of the file:

   ```
   JAVA_OPTS="$JAVA_OPTS -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Djboss.as.management.blocking.timeout=480 -
   Duser.timezone=GMT -Xmx2048m -XX:MaxMetaspaceSize=512m"
   ```

Important: On JDK 11, the setting `-Djava.locale.providers=JRE,COMPAT,CLDR` is required to display four-digit years. Since JDK 9, the Unicode Common Locale Data Repository (CLDR) is the default locales provider. CLDR does not provide years in a four-digit format (see LPS-87191). This setting works around the issue by using JDK 8's default locales provider.

Note: If you plan on using the IBM JDK with your JBoss server, you must complete some additional steps. First, navigate to the `$JBOSS_HOME/modules/com/liferay/portal/main/module.xml` file and insert the following dependency within the `<dependencies>` element:

```
<module name="ibm.jdk" />
```

Then navigate to the `$JBOSS_HOME/modules/system/layers/base/sun/jdk/main/module.xml` file and insert the following path names inside the `<paths>...</paths>` element:

```
<path name="com/sun/crypto" />
<path name="com/sun/crypto/provider" />
<path name="com/sun/image/codec/jpeg" />
<path name="com/sun/org/apache/xml/internal/resolver" />
<path name="com/sun/org/apache/xml/internal/resolver/tools" />
```

The added paths resolve issues with portal deployment exceptions and image uploading problems.

Checkpoint:
At this point, you've finished configuring the application server's JVM settings.

1. The file encoding, user time-zone, preferred protocol stack have been set in the `JAVA_OPTS` in the `standalone.conf.bat` file.

2. The default amount of memory available has been increased.

The prescribed script modifications are now complete for your Liferay DXP installation on JBoss. Next you'll configure the database and mail.

Database Configuration

The easiest way to handle your database configuration is to let Liferay DXP manage your data source. The Basic Configuration page lets you configure Liferay DXP's built-in data source. If you want to use the built-in data source, skip this section.

If you want JBoss to manage your data source, follow these steps:

1. Add the data source inside the `$JBOSS_HOME/standalone/configuration/standalone.xml` file's the `<datasources>` element.

```
<datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-java-
context="true" use-ccm="true">
    <connection-url>[place the URL to your database here]</connection-url>
    <driver>[place the driver name here]</driver>
    <security>
        <user-name>[place your user name here]</user-name>
        <password>[place your password here]</password>
    </security>
</datasource>
```

Make sure to replace the database URL, user name, and password with the appropriate values.

```
**Note:** If you must change your datasource `jndi-name` to something
different, you must also edit the `datasource` element in the
`<default-bindings>` tag.
```

---

2. Add the driver to the `standalone.xml` file's `<drivers>` element also found within the `<datasources>` element.

```
<drivers>
    <driver name="[name of driver must match name above]" module="com.liferay.portal">
        <driver-class>[place your JDBC driver class here]</driver-class>
    </driver>
</drivers>
```

A final data source subsystem that uses MySQL should look like this:

```
<subsystem xmlns="urn:jboss:domain:datasources:5.0">
    <datasources>
        <datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-
java-context="true" use-ccm="true">
            <connection-url>jdbc:mysql://localhost/lportal</connection-url>
            <driver>mysql</driver>
            <security>
```

```
                <user-name>root</user-name>
                <password>root</password>
            </security>
        </datasource>
        <drivers>
            <driver name="mysql" module="com.liferay.portal"/>
        </drivers>
    </datasources>
</subsystem>
```

3. In a `portal-ext.properties` file in your Liferay Home, specify your data source:

```
jdbc.default.jndi.name=java:jboss/datasources/ExampleDS
```

Now that you've configured your data source, the mail session is next.

Mail Configuration

As with database configuration, the easiest way to configure mail is to let Liferay DXP handle your mail session. If you want to use Liferay DXP's built-in mail session, skip this section and configure the mail session in the Control Panel.

If you want to manage your mail session with JBoss, follow these steps:

1. Specify your mail subsystem in the `$WILDFLY_HOME/standalone/configuration/standalone.xml` file like this:

```
<subsystem xmlns="urn:jboss:domain:mail:3.0">
    <mail-session jndi-name="java:jboss/mail/MailSession" >
        <smtp-server ssl="true" outbound-socket-binding-ref="mail-smtp">
            <login username="USERNAME" password="PASSWORD"/>
        </smtp-server>
   </mail-session>
</subsystem>
...
<socket-binding-group name="standard-sockets" default-interface="public" port-offset="${jboss.socket.binding.port-
offset:0}">
...
<outbound-socket-binding name="mail-smtp">
        <remote-destination host="smtp.gmail.com" port="465"/>
    </outbound-socket-binding>
</socket-binding-group>
```

2. In your `portal-ext.properties` file in Liferay Home, reference your mail session:

```
mail.session.jndi.name=java:jboss/mail/MailSession
```

You've got mail! Next, you'll deploy Liferay DXP to your JBoss app server.

Deploying Liferay DXP

1. If the folder `$JBOSS_HOME/standalone/deployments/ROOT.war` already exists in your JBoss installation, delete all of its subfolders and files. Otherwise, create a new folder called `$JBOSS_HOME/standalone/deployments/ROOT.war`.

2. Unzip the Liferay DXP `.war` file into the `ROOT.war` folder.

3. To trigger deployment of `ROOT.war`, create an empty file named `ROOT.war.dodeploy` in your `$JBOSS_HOME/standalone/deployments/` folder. On startup, JBoss detects the presence of this file and deploys it as a web application.

4. Start the JBoss application server by navigating to `$JBOSS_HOME/bin` and running `standalone.bat` or `standalone.sh`.

Congratulations; you've now deployed Liferay DXP on JBoss!

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass gatherExternProperties
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass checkControlFlow
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patter
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
 current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patte
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
```

85.7 Installing Liferay DXP on tc Server

Installing Liferay DXP on tc server requires these *Additional Files* for DXP or Portal CE::

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Liferay DXP requires a Java JDK 8 or 11.

Note: Please see the compatibility matrix for information on supported JDKs, databases, and environments.

Here are the basic steps for installing Liferay DXP on tc Server:

- Installing Liferay DXP dependencies to your application server

- Configuring your application server for Liferay DXP

- Deploying the Liferay DXP WAR file to your application server

---

Note: You'll see the term *Liferay Home* used in this installation guide. *Liferay Home* refers to the folder containing your tc Server instance and some Liferay DXP-specific folders:. `data`, `deploy`, `licenses`, and `osgi` folders.

Creating a tc Server

1. Download and unzip a tc Server, available here. This is referred to as `[TCSERVER_INSTANCE_HOME]`.

2. Create a folder called `servers` inside `[TCSERVER_INSTANCE_HOME]`. (e.g. `/opt/pivotal-tc-server/servers`). This folder becomes *Liferay Home* (see note above), and you should not confuse the two.

3. Next, create an instance called *dxp-server* to deploy Liferay DXP. Navigate to the `[TCSERVER_INSTANCE_HOME]/developer-4.0.0.RELEASE` folder and run this command:

   ```
   tcserver create -i ../servers dxp-server
   ```

   By default, new instances are created in the `[TCSERVER_INSTANCE_HOME]/instances` folder. The `-i` specifies where to create the folder. The instance location can also be changed via a properties file. See the tc Server documentation for more information.

Checkpoint:

1. A new folder called `servers` has been created.

2. A new folder called `dxp-server` has been created inside the `servers` folder. The following folders have been created inside the `dxp-server` folder:

   - `bin`
   - `conf`
   - `lib`
   - `logs`
   - `temp`
   - `webapps`
   - `work`

Next you can install the required dependencies.

Installing Liferay DXP Dependencies

Liferay DXP depends on additional JARs that aren't included with tc Server by default. There are even more JARs that you'd find in a Liferay DXP bundle that are not required but can be useful. If you don't have a Liferay DXP bundle, you can download the required JARs from third parties, as described below.

Note: Many required and useful JARs are pre-installed when you build Liferay DXP from the source code or download a Liferay DXP bundle. If you want to acquire all of the JARs that ship with a Liferay DXP bundle quickly, using one of these sources might save you time.

---

1. Extract the JARs from the dependencies ZIP to the `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/lib` folder.

2. Download the following JARs or copy them from a Liferay DXP bundle to the `[TCSERVER_INSTANCE_HOME]/servers` `server/lib` folder:

   - `activation.jar`
   - `ccpp.jar`
   - `jms.jar`
   - `jta.jar`
   - `jutf7.jar`
   - `mail.jar`
   - `persistence.jar`
   - `support-tomcat.jar`

3. Download a database driver `.jar` file and copy it to the `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/lib` folder. For a list of supported databases, see Liferay's compatibility matrix

4. Create an `osgi` folder in your *Liferay Home*. Extract the folders (i.e., `configs`, `core`, and more) from OSGi ZIP file to the `osgi` folder. The `osgi` folder provides the necessary modules for Liferay DXP's OSGi runtime.

Configuring tc Server

There are a few configuration edits to make so Liferay DXP runs well on tc Server. All of these configuration changes should be made in your tc Server runtime instance.

1. Navigate to the `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/bin` folder. In `setenv.sh` replace this line:

   ```
   JVM_OPTS="-Xmx512M -Xss256K"
   ```

   with this one

   ```
   JVM_OPTS="-Dfile.encoding=UTF-8 -Duser.timezone=GMT -Xmx2048M -Xss512K -XX:MaxMetaspaceSize=512m"
   ```

   In `setenv.bat` replace

   ```
   set JVM_OPTS=-Xmx512M -Xss256K
   ```

   with

   ```
   set JVM_OPTS=-Dfile.encoding=UTF-8 -Duser.timezone=GMT -Xmx2048M -Xss512K -XX:MaxMetaspaceSize=512m
   ```

```
**Important:** For Liferay DXP to work properly, the application server JVM
must use the `GMT` time zone and `UTF-8` file encoding.
```

---

2. Next, you should make sure that UTF-8 URI encoding is used consistently. Open `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/conf/server.xml` and make sure the `Connector` tag includes setting the `URIEncoding` to `UTF-8`.

```
<Connector acceptCount="100"
           connectionTimeout="20000"
           executor="tomcatThreadPool"
           maxKeepAliveRequests="15"
           port="${bio.http.port}"
           protocol="org.apache.coyote.http11.Http11Protocol"
           redirectPort="${bio.https.port}"
           URIEncoding="UTF-8" />
```

3. If you're installing Liferay DXP and tc Server on Windows, open `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/conf/wrapper.conf` and replace

   ```
   wrapper.java.additional.8=-Xss256K
   ```

   with

   ```
   wrapper.java.additional.8=-Xmx2048M
   ```

   and add these properties:

   ```
   wrapper.java.additional.9=-Xss512K
   wrapper.java.additional.10=-XX:MaxMetaspaceSize=256M
   wrapper.java.additional.11=-Dfile.encoding=UTF-8
   ```

```
**Important:** For Liferay DXP to work properly, the application server JVM
must use the `GMT` time zone and `UTF-8` file encoding. If your Java wrapper
doesn't already specify the `GMT` time zone, add an entry for it:

    wrapper.java.additional.12=-Duser.timezone=GMT
```

```
 **Important:** On JDK 11, the setting `-Djava.locale.providers=JRE,COMPAT,CLDR` is required to display four-digit years. Since JDK 9, the Unicode Co
digit format (see [LPS-87191](https://issues.liferay.com/browse/LPS-87191)). This setting works around the issue by using JDK 8's default locales pro
```

---

4. Finally, open `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/conf/web.xml` and add the following configuration after `<load-on-startup>3</load-on-startup>`:

   ```
   <init-param>
       <param-name>compilerSourceVM</param-name>
       <param-value>1.8</param-value>
   </init-param>
   <init-param>
       <param-name>compilerTargetVM</param-name>
       <param-value>1.8</param-value>
   </init-param>
   ```

Database Configuration

The easiest way to handle your database configuration is to let Liferay DXP manage your data source. If you want to use Liferay DXP's built-in data source, you can skip this section.

If you want tc Server to manage your data source, follow these steps:

1. Make sure your database server is installed and working. If it's installed on a different machine, make sure it's accessible from your Liferay DXP machine.

2. Add your data source as a resource in the context of your web application specified in `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/conf/Catalina/localhost/ROOT.xml` (create this file if you don't have it already). An example configuration is shown below:

```
<Context crossContext="true" path="">
  <Resources>
        <PreResources
            base="${catalina.base}/lib/ext/portal"
            className="com.liferay.support.tomcat.webresources.ExtResourceSet"
            webAppMount="/WEB-INF/lib"
        />
    </Resources>
  <Resource
      name="jdbc/LiferayPool"
      auth="Container"
      type="javax.sql.DataSource"
      driverClassName="com.mysql.jdbc.Driver"
      url="jdbc:mysql://localhost/lportal?useUnicode=true&amp;characterEncoding=UTF-8"
      username="root"
      password="root"
      maxActive="100"
      maxIdle="30"
      maxWait="10000"
  />
</Context>
```

   The resource definition above is for a MySQL database named `lportal`that has a user named `root` whose password is `root`. Replace these values with your own.

3. In a `portal-ext.properties` file in Liferay Home, specify your data source:

```
jdbc.default.jndi.name=jdbc/LiferayPool
```

Your data source is now configured. Next set up the mail session.

Mail Configuration

As with database configuration, the easiest way to configure mail is to let Liferay DXP handle your mail session. If you want to use Liferay DXP's built-in mail session, skip this section and configure the mail session in the Control Panel.

If you want to manage your mail session with tc Server, follow these steps:

1. Edit `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/conf/Catalina/localhost/ROOT.xml` and add your mail session as a `Resource` in your web application `Context`. Make sure to replace the example mail session values with your own.

```
<Context crossContext="true" path="">
  <Resources>
    <PreResources
      base="${catalina.base}/lib/ext/portal"
      className="com.liferay.support.tomcat.webresources.ExtResourceSet"
      webAppMount="/WEB-INF/lib"
    />
  </Resources>
  ...
  <Resource
      name="mail/MailSession"
      auth="Container"
      type="javax.mail.Session"
      mail.pop3.host="pop.gmail.com"
      mail.pop3.port="110"
      mail.smtp.host="smtp.gmail.com"
      mail.smtp.port="465"
      mail.smtp.user="user"
      mail.smtp.password="password"
      mail.smtp.auth="true"
      mail.smtp.starttls.enable="true"
      mail.smtp.socketFactory.class="javax.net.ssl.SSLSocketFactory"
      mail.imap.host="imap.gmail.com"
      mail.imap.port="993"
      mail.transport.protocol="smtp"
      mail.store.protocol="imap"
  />
</Context>
```

2. In your `portal-ext.properties` file in Liferay Home, reference your mail session:

```
mail.session.jndi.name=mail/MailSession
```

You created a mail session for tc Server to manage and configured Liferay DXP to use it.

Deploying Liferay DXP

Now you're ready to deploy Liferay DXP using your Liferay DXP WAR file.

1. If you are manually installing Liferay DXP on a clean tc Server instance, delete the contents of the `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/webapps/ROOT` directory. This removes the default home page.

2. Extract the Liferay DXP `.war` file to `[TCSERVER_INSTANCE_HOME]/servers/dxp-server/webapps/ROOT`.

   Now it's time to launch Liferay DXP!

3. Start tc Server by navigating to the runtime instance's `[TCSERVER_INSTANCE_HOME]/developer-4.0.0.RELEASE` folder and run the following command:

```
tcserver run -i ../servers dxp-server
```

Congratulations on successfully installing and deploying Liferay DXP on tc Server!

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass gatherExternProperties
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass checkControlFlow
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patter
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
 current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patte
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread express
```

85.8 Installing Liferay DXP on WebLogic 12c R2

Although you can install Liferay DXP in a WebLogic Admin Server, this isn't recommended. It's a best practice to install web apps, including Liferay DXP, in a WebLogic Managed server. Deploying to a Managed Server lets you start or shut down Liferay DXP more quickly and facilitates transitioning into a cluster configuration. This article therefore focuses on installing Liferay DXP in a Managed Server.

Before getting started, create your Admin and Managed Servers. See WebLogic's documentation for instructions on setting up and configuring Admin and Managed Servers.

You should also read the following articles to familiarize yourself with Liferay DXP's general installation steps:

- Installing Liferay DXP
- Preparing for Install
- Installing Liferay DXP Manually

Before proceeding, you must download these *Additional Files* for DXP or Portal CE:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Liferay DXP requires a Java JDK 8 or 11.

Note: Please see the compatibility matrix for information on supported JDKs, databases, and environments.

Without any further ado, get ready to install Liferay DXP in WebLogic!

Configuring WebLogic's Node Manager

WebLogic's Node Manager starts and stops managed servers.

If you're running WebLogic on a UNIX system other than Solaris or Linux, use the Java Node Manager, instead of the native version of the Node Manager, by configuring these Node Manager properties in the `domains/your_domain_name/nodemanager/nodemanager.properties` file:

```
NativeVersionEnabled=false

StartScriptEnabled=true
```

Note: By default, SSL is used with Node Manager. If you want to disable SSL during development, for example, set `SecureListener=false` in your `nodemanager.properties` file.

See Oracle's Configuring Java Node Manager documentation for details.

Configuring WebLogic

Configure the JVM and other options in a `setUserOverridesLate` WebLogic startup script and in your Managed Server UI.

1. Create a `setUserOverridesLate.sh` script in `[Your Domain]/bin`.

2. Add the following settings.

   ```
   export DERBY_FLAG="false"
   export JAVA_OPTIONS="${JAVA_OPTIONS} -Dfile.encoding=UTF-8 -Duser.timezone=GMT -da:org.apache.lucene... -da:org.aspectj..."
   export JAVA_PROPERTIES="-Dfile.encoding=UTF-8 ${JAVA_PROPERTIES} ${CLUSTER_PROPERTIES}"
   export MW_HOME="[place your WebLogic Server folder path here]"
   export USER_MEM_ARGS="-Xms512m -Xmx2048m"
   export WLS_MEM_ARGS_64BIT="-Xms512m -Xmx2048m"
   export WLS_MEM_ARGS_32BIT="-Xms512m -Xmx2048m"
   ```

~~~~
**Important:** For Liferay DXP to work properly, the application server JVM
must use the `GMT` time zone and `UTF-8` file encoding.
~~~~

~~~~
**Important:** On JDK 11, the setting `-Djava.locale.providers=JRE,COMPAT,CLDR` is required to display four-digit years. Since JDK 9, the Unicode Co
digit format (see [LPS-87191](https://issues.liferay.com/browse/LPS-87191)). This setting works around the issue by using JDK 8's default locales pro
~~~~

~~~~
The `DERBY_FLAG` setting disables the Derby server built in to WebLogic, as DXP does not require this server.

`JAVA_OPTIONS` sets DXP's UTF-8 requirement, Lucene usage, and Aspect Oriented Programming via AspectJ.

`JAVA_PROPERTIES` also sets DXP's UTF-8 requirement.

Set `MW_HOME` to the folder containing the WebLogic server on the machine. For example,

```bash
export MW_HOME="/Users/ray/Oracle/wls12210"
```

The `*_MEM_ARGS` variables set DXP's starting and maximum heap memory capacity.
~~~~

5. You must also ensure that the Node Manager sets Liferay DXP's memory requirements when starting the Managed Server. In the Admin Server's console UI, navigate to the Managed Server you want to deploy Liferay DXP to and select the *Server Start* tab. Enter the following parameters into the *Arguments* field:

   ```
   -Xmx2048m -XX:MaxMetaspaceSize=512m
   ```

   Click *Save* when you're finished.

Next, you'll set some Liferay DXP-specific properties for your Liferay DXP installation.

Setting Liferay DXP Properties

Before installing Liferay DXP, you must set the *Liferay Home* folder's location via the `liferay.home` property in a `portal-ext.properties` file. You can also use this file to override other Liferay DXP properties that you may need.

First, decide which folder you want to serve as Liferay Home. In WebLogic, your domain's folder is generally Liferay Home, but you can choose any folder on your machine. Then create your `portal-ext.properties` file and add the `liferay.home` property:

```
liferay.home=/full/path/to/your/liferay/home/folder
```

Remember to change this file path to the location on your machine that you want to serve as Liferay Home.

Now that you've created your `portal-ext.properties` file, you must put it inside the Liferay DXP WAR file. Expand the Liferay DXP WAR file and place `portal-ext.properties` in the `WEB-INF/classes` folder. Later, you can deploy the expanded archive to WebLogic. Alternatively, you can re-WAR the expanded archive for later deployment. In either case, Liferay DXP reads your property settings once it starts up.

If you need to make any changes to `portal-ext.properties` after Liferay DXP deploys, you can find it in your domain's `autodeploy/ROOT/WEB-INF/classes` folder. Note that the `autodeploy/ROOT` folder contains the Liferay DXP deployment.

Next, you'll install Liferay DXP's dependencies.

Installing Liferay DXP Dependencies

You must now install Liferay DXP's dependencies. Recall that earlier you downloaded two ZIP files containing these dependencies. Install their contents now:

1. `liferay-dxp-dependencies-[version].zip`: Unzip this file and place its contents in your WebLogic domain's `lib` folder.

2. `liferay-dxp-osgi-[version].zip`: Unzip this file and place its contents in the `[Liferay Home]/osgi` folder (create this folder if it doesn't exist).

3. Download your database driver `.jar` file and copy it to your domain's `lib` folder. Please see the compatibility matrix for a list of supported databases.

Next, you'll configure your database.

Database Configuration

Use the following procedure if you want WebLogic to manage your database for Liferay DXP. You can skip this section if you want to use Liferay DXP's built-in Hypersonic database.

1. Log in to your AdminServer console.

2. In the *Domain Structure* tree, find your domain and navigate to *Services* → *JDBC* → *Data Sources*.

3. To create a new data source, click *New*. Fill in the *Name* field with `Liferay Data Source` and the *JNDI Name* field with `jdbc/LiferayPool`. Select your database type and driver. For example, MySQL is *MySQL's Driver (Type 4) Versions:using com.mysql.jdbc.Driver*. Click *Next* to continue.

4. Accept the default settings on this page and click *Next* to move on.
5. Fill in your database information for your MySQL database.
6. If using MySQL, add the text `?useUnicode=true&characterEncoding=UTF-8&\useFastDateParsing=false` to the URL line and test the connection. If it works, click *Next*.
7. Select the target for the data source and click *Finish*.
8. You must now tell Liferay DXP about the JDBC data source. Create a `portal-ext.propreties` file in your Liferay Home directory, and add the line `jdbc.default.jndi.name=jdbc/LiferayPool`.

Alternatively, you can make the above configuration strictly via properties in the `portal-ext.properties` file. To do so, place the following properties and values in the file. Be sure to change the `your*` values with the values appropriate for your database's configuration (if using MySQL):

```
jdbc.default.driverClassName=com.mysql.jdbc.Driver
jdbc.default.url=jdbc:mysql://your.db.ip.address/yourdbname?useUnicode?useUnicode=true&characterEncoding=UTF-8&useFastDateParsing=false
jdbc.default.username=yourdbuser
jdbc.default.password=yourdbpassword
```

Next, you'll configure your mail session.

Mail Configuration

If you want WebLogic to manage your mail session, use the following procedure. If you want to use Liferay's built-in mail session (recommended), skip this section.

1. Start WebLogic and log in to your Admin Server's console.
2. Select *Services → Mail Sessions* from the *Domain Structure* box on the left hand side of your Admin Server's console UI.
3. Click *New* to begin creating a new mail session.
4. Name the session *LiferayMail* and give it the JNDI name `mail/MailSession`. Then fill out the *Session Username*, *Session Password*, *Confirm Session Password*, and *JavaMail Properties* fields as necessary for your mail server. See the WebLogic documentation for more information on these fields. Click *Next* when you're done.
5. Choose the Managed Server that you'll install Liferay DXP on, and click *Finish*. Then shut down your Managed and Admin Servers.
6. With your Managed and Admin servers shut down, add the following property to your `portal-ext.properties` file in Liferay Home:

   ```
   mail.session.jndi.name=mail/MailSession
   ```

   Liferay DXP references your WebLogic mail session via this property setting. If you've already deployed Liferay DXP, you can find your `portal-ext.properties` file in your domain's `autodeploy/ROOT/WEB-INF/classes` folder.

Your changes take effect upon restarting your Managed and Admin servers.

Deploying Liferay DXP

As mentioned earlier, although you can deploy Liferay DXP to a WebLogic Admin Server, you should instead deploy it to a WebLogic Managed Server. Dedicating the Admin Server to managing other servers that run your apps is a best practice.

Follow these steps to deploy Liferay DXP to a Managed Server:

1. Make sure the Managed Server you want to deploy Liferay DXP to is shut down.
2. In your Admin Server's console UI, select *Deployments* from the *Domain Structure* box on the left hand side. Then click *Install* to start a new deployment.
3. Select the Liferay DXP WAR file or its expanded contents on your file system. Alternatively, you can upload the WAR file by clicking the *Upload your file(s)* link. Click *Next*.
4. Select *Install this deployment as an application* and click *Next*.
5. Select the Managed Server you want to deploy Liferay DXP to and click *Next*.
6. If the default name is appropriate for your installation, keep it. Otherwise, give it a name of your choosing and click *Next*.
7. Click *Finish*. After the deployment finishes, click *Save* if you want to save the configuration.
8. Start the Managed Server where you deployed Liferay DXP. Liferay DXP precompiles all the JSPs and then launches.

Nice work! Now you're running Liferay DXP on WebLogic.

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass gatherExternProperties
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass checkControlFlow
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patter
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
 current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patte
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
```

85.9 Installing Liferay DXP on WebSphere

IBM ® WebSphere ® is a trademark of International Business Machines Corporation, registered in many jurisdictions worldwide.

Tip: Throughout this installation and configuration process, WebSphere prompts you to click *Save* to apply changes to the Master Configuration. Do so intermittently to save your changes.

For Liferay DXP to work correctly, WebSphere 8.5.5 Fix Pack 11 (or later) or WebSphere 9 (Fix Pack 7 is the latest) must be installed. You can find more information about this fix pack here.

Please also note that Liferay DXP doesn't support the WebSphere Application Liberty Profile.

You should also read the following articles to familiarize yourself with Liferay DXP's general installation steps:

- Preparing for Install
- Installing Liferay DXP

Note that the *Liferay Home* folder is important to the operation of Liferay DXP. In Liferay Home, Liferay DXP creates certain files and folders that it needs to run. On WebSphere, Liferay Home is typically `[Install Location]/WebSphere/AppServer/profiles/your-profile/liferay`.

Before proceeding, you must download these *Additional Files* for DXP or Portal CE::

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Liferay DXP requires a Java JDK 8 or 11.

Note: Please see the compatibility matrix for information on supported JDKs, databases, and environments.

Without any further ado, get ready to install Liferay DXP in WebSphere!

Preparing WebSphere for Liferay DXP

When the application server binaries have been installed, start the *Profile Management Tool* to create a profile appropriate for Liferay DXP.

1. Click *Create…*, choose *Application Server*, and then click *Next*.

2. Click the *Advanced* profile creation option and then click *Next*. You need the advanced profile to specify your own values for settings such as the location of the profile and names of the profile, node and host, to assign your own ports, or to optionally choose whether to deploy the administrative console and sample application and also add web-server definitions for IBM HTTP Server. See the WebSphere documentation for more information about these options.

3. Check the box *Deploy the administrative console*. This gives you a web-based UI for working with your application server. Skip the default applications. You'd only install these on a development machine. Click *Next*.

4. Set the profile name and location. Ensure you specify a performance tuning setting other than *Development*, since you're installing a production server. See the WebSphere documentation for more information about performance tuning settings. Click *Next*.

5. Choose node, server, and host names for your server. These are specific to your environment. Click *Next*.

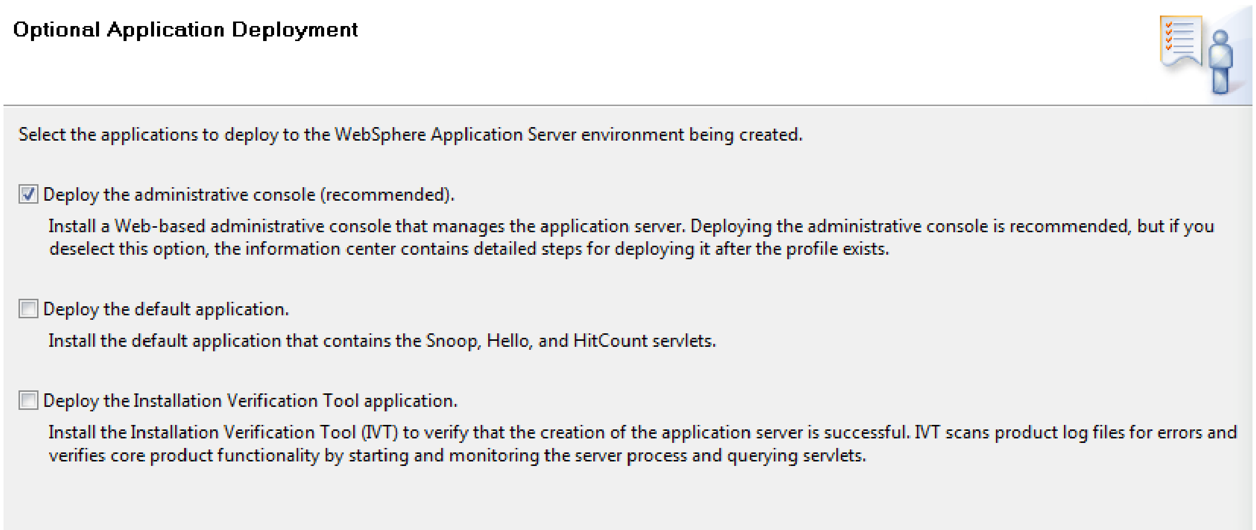

Figure 85.2: Choose the Advanced profile option to specify your own settings.

6. Administrative security in WebSphere is a way to restrict who has access to the administrative tools. You may want to have it enabled in your environment so that a user name and password are required to administer the WebSphere server. See WebSphere's documentation for more information. Click *Next*.

7. Each profile needs a security certificate, which comes next in the wizard. If you don't have certificates already, choose the option to generate a personal certificate and a signing certificate and click *Next*.

8. Once the certificates are generated, set a password for your keystore. Click *Next*.

9. Now you can customize the ports this server profile uses. Be sure to choose ports that are open on your machine. When choosing ports, the wizard detects existing WebSphere installations and if it finds activity, it increments ports by one.

10. Choose whether you want this profile started when the machine starts. Click *Next*.

11. WebSphere ships with IBM HTTP Server, which is a re-branded version of Apache. Choose whether you want a web server definition, so that this JVM receives requests forwarded from the HTTP server. See WebSphere's documentation for details on this. When finished, click *Next*.

12. The wizard then shows you a summary of what you selected, enabling you to keep your choices or go back and change something. When you're satisfied, click *Next*.

WebSphere then creates your profile and finishes with a message telling you the profile was created successfully. Awesome! Your profile is complete. Now there are a few things you need to configure in your application server.

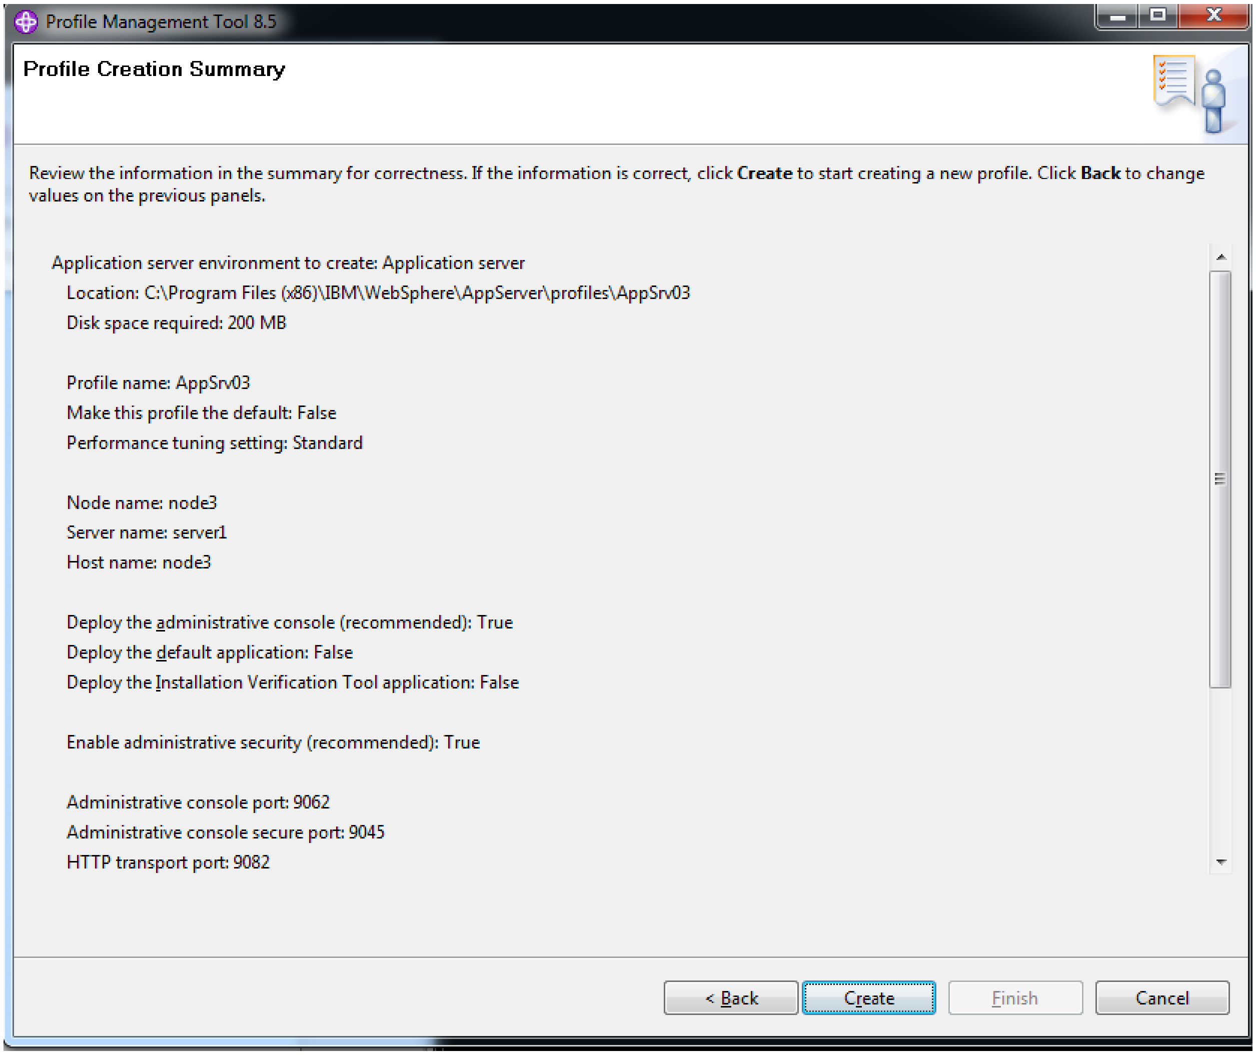

Figure 85.3: Example of the settings before creating the profile.

Configuring the WebSphere Application Server

In this version of WebSphere, servlet filters are not initialized on web application startup, but rather, on first access. This can cause problems when deploying certain apps to Liferay DXP. To configure servlet filters to initialize on application startup (i.e., deployment), set the following `webcontainer` properties in your WebSphere application server:

```
com.ibm.ws.webcontainer.initFilterBeforeInitServlet = true
com.ibm.ws.webcontainer.invokeFilterInitAtStartup = true
```

To set `webcontainer` properties in the WebSphere application server, follow the instructions here in WebSphere's documentation.

Setting up JVM Parameters for Liferay DXP

Next, in the WebSphere profile you created for Liferay DXP, you must set an argument that supports Liferay DXP's Java memory requirements. You'll modify this file:

```
[Install Location]/WebSphere/AppServer/profiles/your-profile/config/cells/your-cell/nodes/your-node/servers/your-server/server.xml
```

Add `maximumHeapSize="2048"` inside the `jvmEntries` tag. For example:

```
<jvmEntries xmi:id="JavaVirtualMachine_1183122130078" ... maximumHeapSize="2048">
```

Note: The JVM parameters used here are defaults intended for initial deployment of production systems. Administrators should change the settings to values that best address their specific environments. These must be tuned depending on need.

Administrators can set the UTF-8 properties in the `server.xml` file. This is required or else special characters will not be parsed correctly. Add the following inside the `jvmEntries` tag:

```
<jvmEntries xmi:id="JavaVirtualMachine_1183122130078" ...genericJvmArguments="-Dfile.encoding=UTF-8 -Duser.timezone=GMT">
```

Important: For Liferay DXP to work properly, the application server JVM must use the `GMT` time zone and `UTF-8` file encoding.

Alternately, you can set the UTF-8 properties from the WebSphere Admin Console. (See below.)

Important: On JDK 11, the setting `-Djava.locale.providers=JRE,COMPAT,CLDR` is required to display four-digit years. Since JDK 9, the Unicode Common Locale Data Repository (CLDR) is the default locales provider. CLDR does not provide years in a four-digit format (see LPS-87191). This setting works around the issue by using JDK 8's default locales provider.

Removing the secureSessionCookie Tag

In the same profile, you should delete a problematic `secureSessionCookie` tag that can cause Liferay DXP startup errors. Note that this is just a default setting; once Liferay DXP is installed, you should tune it appropriately based on your usage.

In `[Install Location]/WebSphere/AppServer/profiles/your-profile/config/cells/your-cell/cell.xml`, Delete the `secureSessionCookie` tag containing `xmi:id="SecureSessionCookie_1"`.

If this tag is not removed, an error similar to this may occur:

```
WSVR0501E: Error creating component com.ibm.ws.runtime.component.CompositionUnitMgrImpl@d74fa901
com.ibm.ws.exception.RuntimeWarning: com.ibm.ws.webcontainer.exception.WebAppNotLoadedException: Failed to load webapp: Failed to load webapp: SRVE8
```

Installing Liferay DXP's Dependencies

You must now install Liferay DXP's dependencies. Recall that earlier you downloaded two ZIP files containing these dependencies. Install their contents now:

1. `liferay-dxp-dependencies-[version].zip`: Unzip this file and place its contents in your WebSphere application server's `[Install Location]/WebSphere/AppServer/lib/ext` folder. If you have a JDBC database driver JAR, copy it to this location as well.

2. `liferay-dxp-osgi-[version].zip`: Unzip this file and place its contents in the `[Liferay Home]/osgi` folder (create this folder if it doesn't exist). This is typically `[Install Location]/WebSphere/AppServer/profiles/your-profile/liferay/osgi`.

Installing the DXP portlet.jar

DXP's `portlet.jar` (version 3) is backwards-compatible. It is included with the Dependencies ZIP that you unzipped above. WebSphere contains an older `portlet.jar` version which must be overridden.

1. In your `[Install Location]/WebSphere/AppServer/profiles/your-profile/` folder, create a folder called `app_shared_libraries`.

2. Move DXP's `portlet.jar` from the `[Install Location]/WebSphere/AppServer/lib/ext` folder to the `app_shared_libraries` folder you created.

3. Follow IBM's steps for using a server associated shared library; make sure to choose *Classes loaded with local class loader first (parent_Last)* on step 4d.

4. Save the configuration.

Ensuring That the DXP Portlet.jar is Loaded First

In addition to placing DXP's `portlet.jar` in a server associated shared library, configure the `config.ini` file so that it is loaded first.

1. Open the `[Install Location]/WebSphere/AppServer/configuration/config.ini` file.
2. Find the property `com.ibm.CORBA,com.ibm`.
3. Insert the property `javax.portlet,javax.portlet.filter,javax.portlet.annotations` after `com.ibm.CORBA` and before `com.ibm`.
4. Save the file.

Start the server profile you created for Liferay DXP. Once it starts, you're ready to configure your database.

Database Configuration

If you want WebSphere to manage the database connections, follow the instructions below. Note this is not necessary if you plan to use Liferay DXP's standard database configuration; in that case, skip this section. You'll set your database information in Liferay DXP's setup wizard after the install.

Note: Although Liferay DXP's embedded database is fine for testing purposes, you **should not** use it for production Liferay DXP instances.

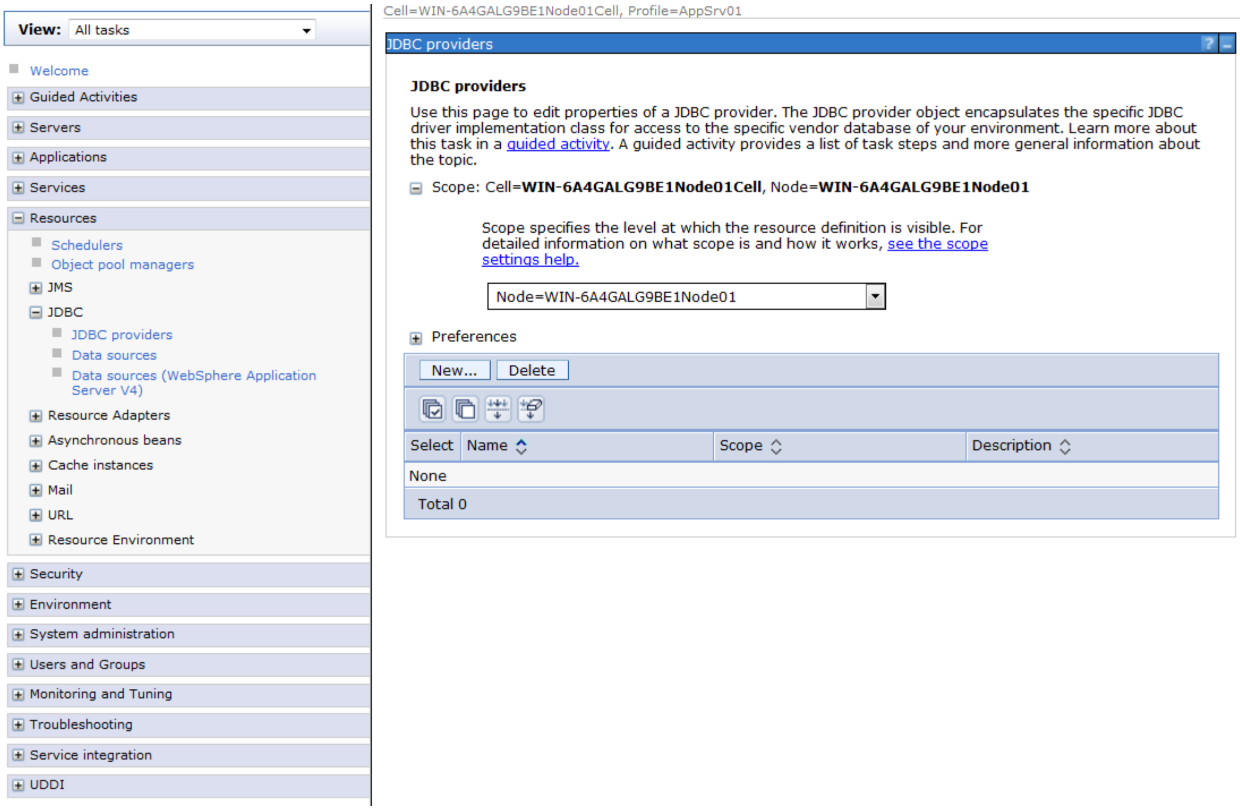

Figure 85.4: WebSphere JDBC providers

1. Start WebSphere.

2. Open the Administrative Console and log in.

3. Click *Resources* → *JDBC Providers*.

4. Select a scope and then click *New*.

5. Select your database type, provider type, and implementation type. If you select a predefined database, the wizard fills in the name and description fields for you. If the database you want to use isn't listed, select *User-defined* from the *Database type* field and then fill in the *Implementation Class Name*. For example, if you use MySQL, select *Database type* → *User-defined*, and then enter `com.mysql.jdbc.jdbc2.optional.MysqlConnectionPoolDataSource` in *Implementation Class Name*. Click *Next* when you are finished.

6. Clear any text in the classpath settings. You already copied the necessary JARs to a location on the server's classpath. Click *Next*.

7. Review your settings and click *Finish*. The final configuration should look like this:

8. Click your new provider configuration when it appears in the table, and then click *Data Sources* under *Additional Properties*. Click *New*.

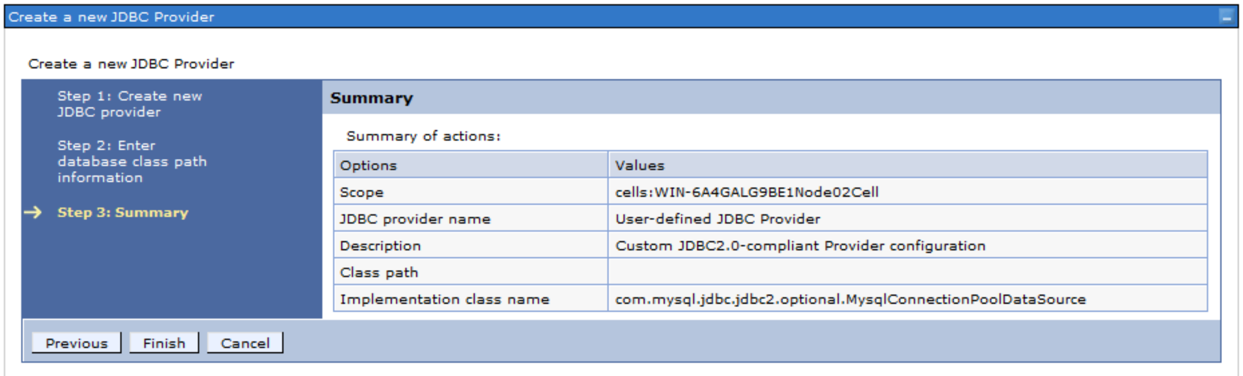

Figure 85.5: Completed JDBC provider configurations.

9. Enter *liferaydatabasesource* in the *Data source name* field and `jdbc/LiferayPool` in the *JNDI name* field. Click *Next*.

10. Click *Next* in the remaining screens of the wizard to accept the default values. Then review your changes and click *Finish*.

11. Click the data source when it appears in the table and then click *Custom Properties*. Now click the *Show Filter Function* button. This is the second from last of the small icons under the *New* and *Delete* buttons.

12. Type *user* into the search terms and click *Go*.

13. Select the *user* property and give it the value of the user name to your database. Click *OK* and save to master configuration.

14. Do another filter search for the *url* property. Give this property a value that points to your database. For example, a MySQL URL would look like this:

    ```
    jdbc:mysql://localhost/lportal?useUnicode=true&characterEncoding=UTF-8&useFastDateParsing=false
    ```

    Click *OK* and save to master configuration.

15. Do another filter search for the *password* property. Enter the password for the user ID you added earlier as the value for this property. Click *OK* and save to master configuration.

16. Go back to the data source page by clicking it in the breadcrumb trail. Click the *Test Connection* button. It should connect successfully.

Once you've set up your database, you can set up your mail session.

Mail Configuration

If you want WebSphere to manage your mail sessions, use the following procedure. If you want to use Liferay DXP's built-in mail sessions, you can skip this section.

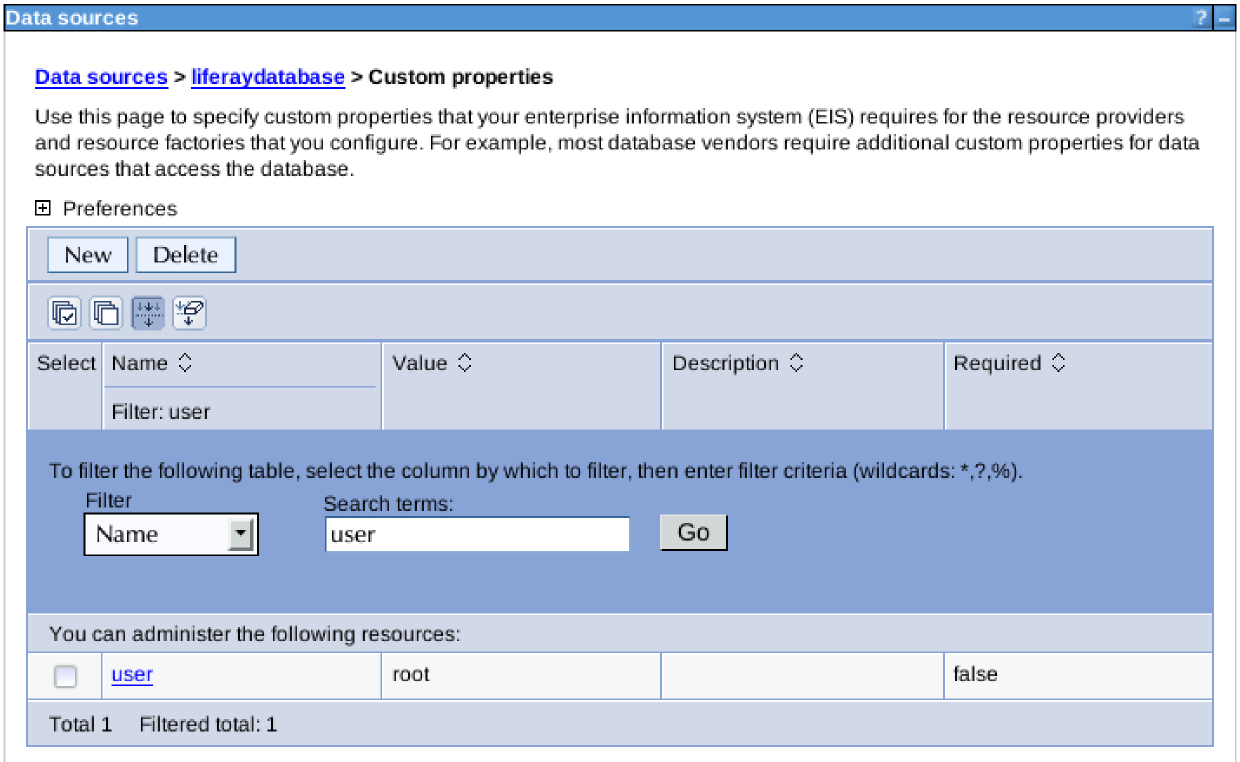

Figure 85.6: Modifying data source properties in WebSphere

Creating a WebSphere-Managed Mail Session (Optional)

1. Click *Resources* → *Mail* → *Mail Providers*.

2. Click the Built-In Mail Provider for your node and server.

3. Click *Mail Sessions* and then click the *New* button.

4. Give your mail session a name of *liferaymail* and a JNDI name of `mail/MailSession`. Fill in the correct information for your mail server in the sections *Outgoing Mail Properties* and *Incoming Mail Properties*. Click *OK* and then save to the master configuration.

5. Click your mail session when it appears in the table and select *Custom Properties* under the *Additional Properties* section. Set any other JavaMail properties required by your mail server, such as the protocol, ports, whether to use SSL, and so on.

6. Click *Security* → *Global Security* and de-select *Use Java 2 security to restrict application access to local resources* if it is selected. Click *Apply*.

Note that you may also need to retrieve a SSL certificate from your mail server and add it to WebSphere's trust store. See WebSphere's documentation for instructions on this.

Verifying WebSphere Mail Provider

To validate that the mail session has been configured correctly, there are a number of ways to test this once the WAR has been deployed, the server has started, and the user has signed in as the system administrator. One quick way to validate is to create a new user with a valid email account. The newly created user should receive an email notification. The logs should display that the SMTP server has been pinged with the correct port number listed.

Enable Cookies for HTTP Sessions

WebSphere restricts cookies to HTTPS sessions by default. If you're using HTTP instead, this prevents users from signing in to Liferay DXP and displays the following error in the console:

```
20:07:14,021 WARN  [WebContainer : 1][SecurityPortletContainerWrapper:341]
User 0 is not allowed to access URL http://localhost:9081/web/guest/home and
portlet com_liferay_login_web_portlet_LoginPortlet
```

This occurs because Liferay DXP can't use the HTTPS cookie when you use HTTP. The end result is that new sessions are created on each page refresh. Follow these steps to resolve this issue in WebSphere:

1. Click *Application Servers* → *server1* → *Session Management* → Enable Cookies
2. De-select *Restrict cookies to HTTPS sessions*
3. Click *Apply*
4. Click *Save*

Enable UTF-8

If you did not add the `-Dfile.encoding=UTF-8` property in the `server.xml`, you can do so in the Administrative Console.

1. Click *Application Servers* → *server1* → *Process definition*.
2. Click *Java Virtual Machine* under *Additional Properties*.
3. Enter `-Dfile.encoding=UTF-8` in the *Generic JVM arguments* field.
4. Click *Apply* and then *Save* to master configuration.

Once the changes have been saved, Liferay DXP can parse special characters if there is localized content.

Deploy Liferay DXP

Now you're ready to deploy Liferay DXP!

1. In WebSphere's administrative console, click *Applications* → *New Application* → *New Enterprise Application*.
2. Browse to the Liferay DXP `.war` file, select it, and click *Next*.

3. Leave *Fast Path* selected and click *Next*. Ensure that *Distribute Application* has been checked and click *Next* again.

4. Choose the WebSphere runtimes and/or clusters where you want Liferay DXP deployed. Click *Next*.

5. Select the virtual host to deploy Liferay DXP on and click *Next*.

6. Map Liferay DXP to the root context (/) and click *Next*.

7. Select the *metadata-complete attribute* setting that you want to use and click *Next*.

8. Ensure that you have made all the correct choices and click *Finish*. When Liferay DXP has installed, click *Save to Master Configuration*.

You've now installed Liferay DXP!

Setting the JDK Version for Compiling JSPs

Liferay DXP requires that its JSPs are compiled to the Java 8 bytecode format. To ensure that WebSphere does this, shut down WebSphere after you've deployed the Liferay DXP `.war` file. Navigate to the `WEB_INF` folder and add the following setting to the `ibm-web-ext.xml` or in most cases the `ibm-web-ext.xmi` file:

```
<jsp-attribute name="jdkSourceLevel" value="18" />
```

The exact path to the `ibm-web-ext.xmi` file depends on your WebSphere installation location and Liferay DXP version, but here's an example:

```
/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/config/cells/localhostNode01Cell/applications/liferayXX.ear/deployments/liferayXX/liferayXX.war/WEB-
INF/ibm-web-ext.xmi
```

Note that the Liferay DXP `.war` comes pre-packaged with the `ibm-web-ext.xmi` file; this format is functionally the same as `.xml` and WebSphere recognizes both formats. For more general information on how WebSphere compiles JSPs see IBM's official documentation for WebSphere Application Server 8.5.5.x or WebSphere Application Server 9.0.0.x.

Now restart WebSphere.

Start Liferay DXP

1. If you plan to use Liferay DXP's setup wizard, skip to the next step. If you wish to use WebSphere's data source and mail session, create a file called `portal-ext.properties` in your Liferay Home folder. Place the following configuration in the file:

   ```
   jdbc.default.jndi.name=jdbc/LiferayPool
   mail.session.jndi.name=mail/MailSession
   setup.wizard.enabled=false
   ```

2. In the WebSphere administrative console, navigate to *Enterprise Applications*, select the Liferay DXP application, and click *Start*. While Liferay DXP is starting, WebSphere displays a spinning graphic. Don't watch it too closely, or you might get hypnotized.

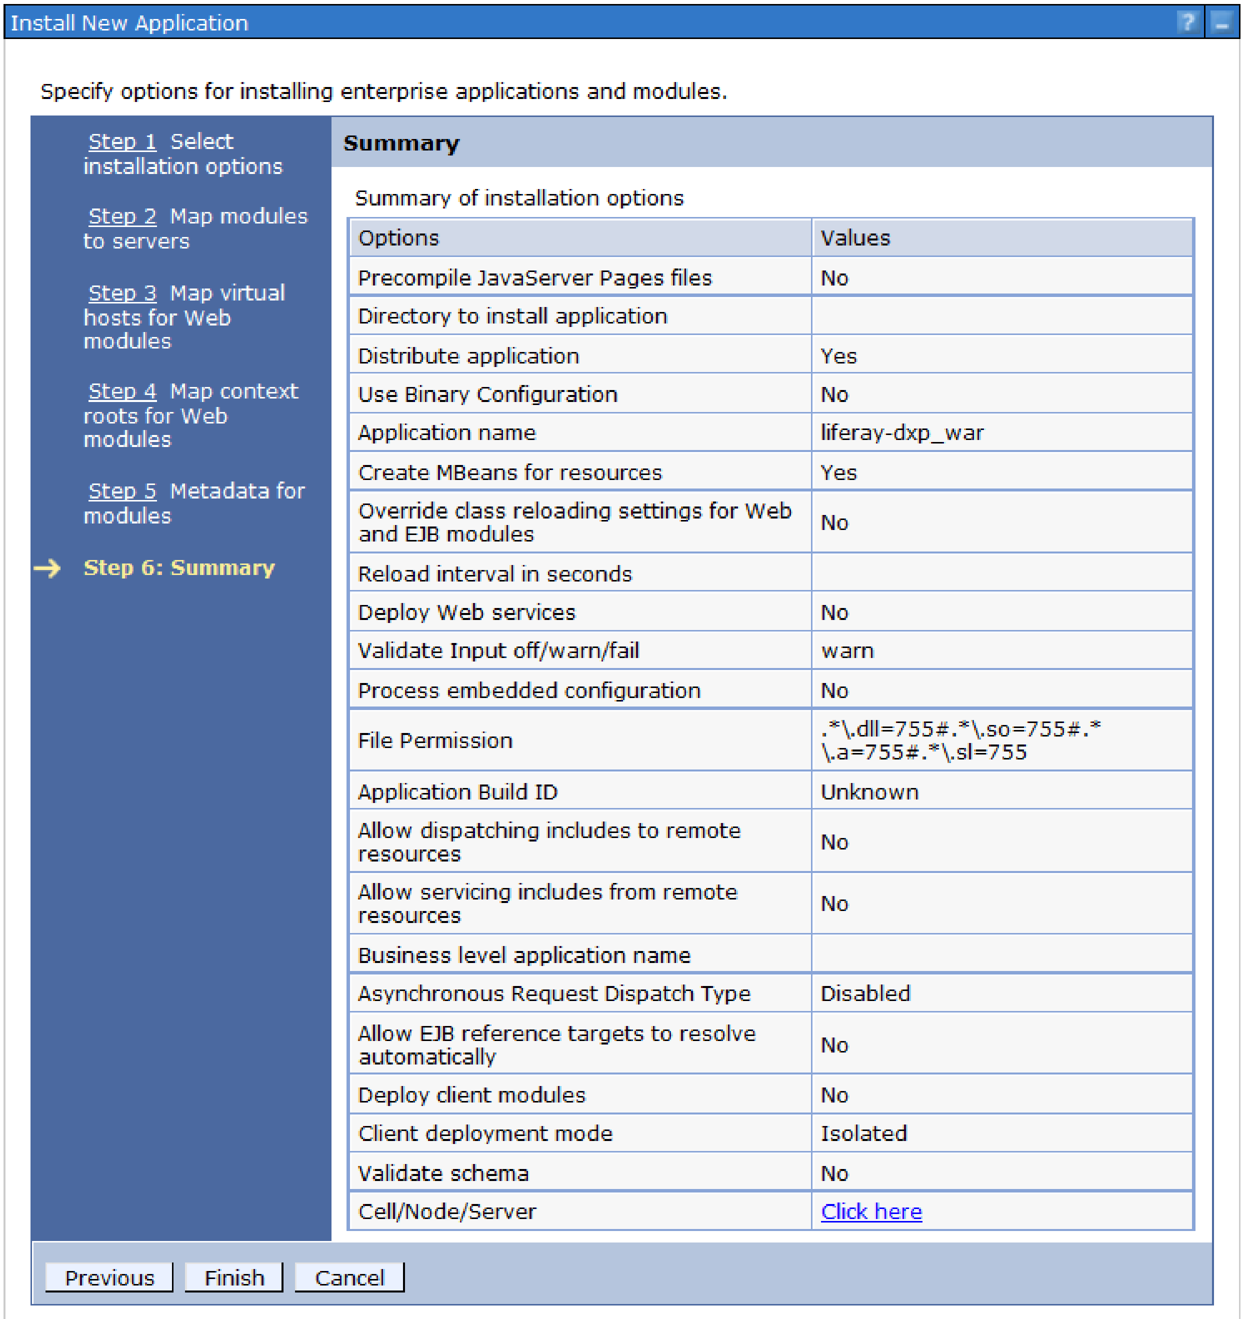

Figure 85.7: Review your deployment options before deploying.

3. In Liferay DXP's setup wizard, select and configure your database type. Click *Finish* when you're done. Liferay DXP then creates the tables it needs in the database.

Congratulations! You've installed Liferay DXP on WebSphere!

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass gatherExternProperties
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 WARNING: Skipping pass checkControlFlow
 May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
 INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patter
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
 current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array patte
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressi
```

85.10 Installing Elasticsearch

Liferay DXP uses Elasticsearch to index its content. By default, it's installed as an embedded service. It works, but it's not a supported configuration for a production server. Feel free to use it while testing or developing, but when you're ready to put your site in production, you must run Elasticsearch as a standalone process. This is better anyway, because it frees you to design your infrastructure the way you want it. If you've got hardware or a VM to spare, you can separate your search infrastructure from Liferay DXP and reap some performance gains by putting search on a separate box. If you're more budget-conscious, you can still increase performance by running Elasticsearch in a separate, individually tunable JVM on the same box.

Before installing Elasticsearch, refer to Preparing to Install Elasticsearch for guidance on configuring the servers to support an Elasticsearch deployment properly.

Installing Elasticsearch is pretty easy and takes only six steps:

1. Find the latest compatible version of Elasticsearch, and then download that version from Elastic's website.

2. Install Elasticsearch by extracting its archive to the system where you want it to run.

3. Install some required Elasticsearch plugins.

4. Name your Elasticsearch cluster.

5. Configure Liferay DXP to connect to your Elasticsearch cluster.

```
**Note:** If you are installing an Elasticsearch version that requires
installation of a new connector application, you must disable several
version-specific modules. The instructions in [Installing Elasticsearch
7](/docs/7-1/deploy/-/knowledge_base/d/installing-elasticsearch-7)
demonstrate this.
```

6. Restart Liferay DXP and reindex your search indexes.

Note: Before continuing, make sure you have set the `JAVA_HOME` environment variable.

If you have multiple JDKs installed, make sure Elasticsearch and Liferay DXP are using the same version and distribution (e.g., Oracle Open JDK 1.8.0_201). You can specify this in `[Elasticsearch Home]/bin/elasticsearch.in.sh`:

```
JAVA_HOME=/path/to/java
```

Consult the Elasticsearch compatibility matrix and the Liferay DXP compatibility matrix to learn more about supported JDK distributions and versions.

Now you'll perform these steps, and when you're done, you'll have a production-ready instance of Liferay DXP up and running.

After you're done following the installation guide, refer to the Configuring Elasticsearch article for more details on configuring Liferay DXP for Elasticsearch. For more information on installing a search engine, see here.

Step One: Find the Right Version of Elasticsearch

To install the latest compatible Elasticsearch version, refer to the Search Engine Compatibility Matrix, and then download that version from Elastic's website.

To inspect the version of the embedded Elasticsearrch server, if Liferay DXP isn't running, start it.

Visit port 9200 on localhost to access the embedded Elasticsearch:

```
http://localhost:9200
```

A JSON document is returned that looks similar to this:

```
{
  "name" : "g0m223N",
  "cluster_name" : "LiferayElasticsearchCluster",
  "cluster_uuid" : "Ii6STs04Tg-XzTVV5h7M2Q",
  "version" : {
    "number" : "6.5.1",
    "build_hash" : "af51318",
    "build_date" : "2018-01-26T18:22:55.523Z",
    "build_snapshot" : false,
    "lucene_version" : "7.1.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "tagline" : "You Know, for Search"
}
```

The version of Elasticsearch that's running is the value of the `"number"` field. In this example, it's 6.5.1.

Shut down the Liferay DXP server. In a local, single-machine testing environment, if you continue without shutting down, the Elasticsearch server you're about to install and start throws errors in the log if its cluster name and HTTP port match the already-running embedded Elasticsearch server. An alternative to shutting down Liferay DXP is to use a different cluster name (i.e., not `LiferayElasticsearchCluster`) and HTTP port (i.e., not `9200`) in the remote Elasticsearch server.

Now that you know the version of Elasticsearch you need, go to Elastic's website and download that version.

Step Two: Install Elasticsearch

Most of this step entails deciding where you want to run Elasticsearch. Do you want to run it on the same machine as Liferay DXP, or do you want to run it on its own hardware? The answer to this question comes down to a combination of the resources you have available and the size of your installation. Regardless of what you decide, either way you get the benefit of a separately tunable search infrastructure.

Once you have a copy of the right version of Elasticsearch, extract it to a folder on the machine where you want it running. That's it!

Step Three: Install Elasticsearch Plugins

Install the following required Elasticsearch plugins:

- `analysis-icu`
- `analysis-kuromoji`
- `analysis-smartcn`
- `analysis-stempel`

To install these plugins, navigate to Elasticsearch Home and enter

```
./bin/elasticsearch-plugin install [plugin-name]
```

Replace *[plugin-name]* with the Elasticsearch plugin's name.

Step Four: Name Your Elasticsearch Cluster

A *cluster* in Elasticsearch is a collection of nodes (servers) identified as a cluster by a shared cluster name. The nodes work together to share data and workload. A one node cluster is discussed here; to create a multi-node cluster, please refer to Elastic's documentation.

Now that you've installed Elastic, it sits in a folder on your machine, which is referred to here as `[Elasticsearch Home]`. To name your cluster, you'll define the cluster name in both Elasticsearch and in Liferay DXP. First, define it in Elasticsearch. Edit the following file:

```
[Elasticsearch Home]/config/elasticsearch.yml
```

Uncomment the line that begins with `cluster.name`. Set the cluster name to whatever you want to name your cluster:

```
cluster.name: LiferayElasticsearchCluster
```

Of course, this isn't a very imaginative name; you may choose to name your cluster `finders_keepers` or something else you can remember more easily. Save the file.

Now you can start Elasticsearch. Run the executable for your operating system from the `[Elasticsearch Home]/bin` folder:

```
./elasticsearch
```

Elasticsearch starts, and one of its status messages includes a transport address:

```
[2018-04-03T15:34:19,784][INFO ][o.e.t.TransportService   ] [g0m223N] publish_address {127.0.0.1:9300}, bound_addresses {[::1]:9300}, {127.0.0.1:930
```

Take note of this address; you'll need to give it to your Liferay DXP server so it can find Elasticsearch on the network.

Step Five: Configure Liferay DXP to Connect to your Elasticsearch Cluster

Now that you're ready to configure Liferay DXP, start it if you haven't already, log in, and then click on *Control Panel* → *Configuration* → *System Settings* → *Search*. Enter the term *elasticsearch* in the search bar and click the *Elasticsearch 6* entry from the list of settings. Now you can configure it. Here are the configuration options to change:

Cluster Name: Enter the name of the cluster as you defined it in Elasticsearch.

Operation Mode: Defaults to EMBEDDED. Change it to REMOTE to connect to a standalone Elasticsearch.

Transport Addresses: Enter a delimited list of transport addresses for Elasticsearch nodes. Here, you'll enter the transport address from the Elasticsearch server you started. The default value is `localhost:9300`, which will work.

When finished, click *Save*. You're almost done.

Step Six: Restart Liferay DXP and Reindex

Stop and restart Liferay DXP. When it's back up, log in as an administrative user and click on *Control Panel* → *Configuration* → *Search* and click the *Execute* button for *Reindex all search indexes*. When you do that, you should see some messages scroll up in the Elasticsearch log.

For more details refer to the Elasticsearch installation guide.

85.11 Trial Plugin Installation

For Liferay customers who are evaluating Liferay DXP on a trial basis, **the plugins can be accessed from within the *Apps* → *Store* (i.e., Marketplace) section of the Control Panel in your product installation**.

Installation Process

Follow the steps below to install a trial plugin:

1. Register a `liferay.com` account (LRDC) account by visiting Liferay's home page (if necessary). Do this by clicking *Sign In/Create Account* button from the top right Profile button.

2. Start your Liferay DXP instance (trial license is OK).

3. After signing in as an Admin in your Liferay DXP trial server, go to the Control Panel → *Apps* → *Store* and configure the `liferay.com` (LRDC) account to log in to the Marketplace. Authorize Marketplace to access your local account.

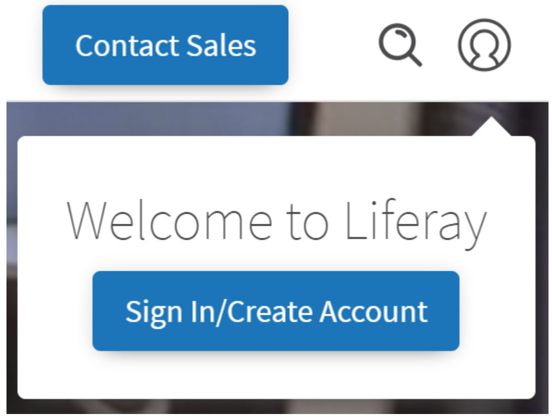

Figure 85.8: Hover over the Profile button and click *Sign In/Create Account*.

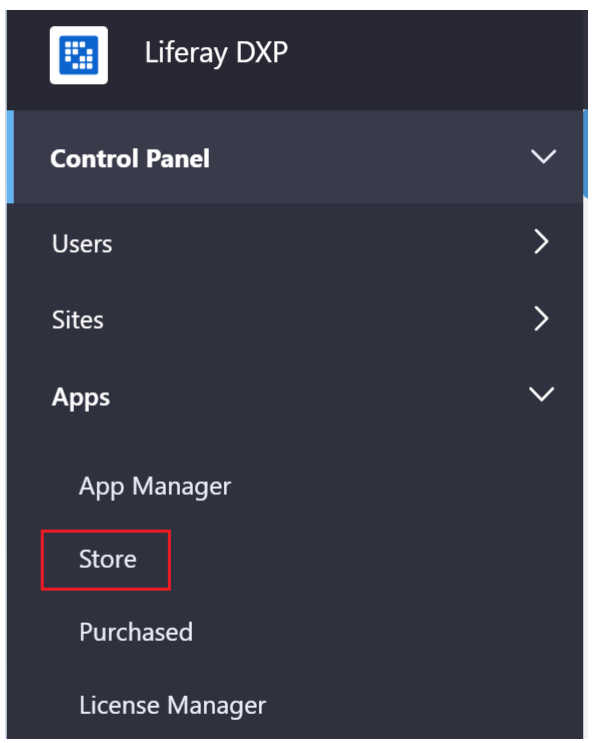

Figure 85.9: Click the *Store* link and authorize Marketplace to access your local account.

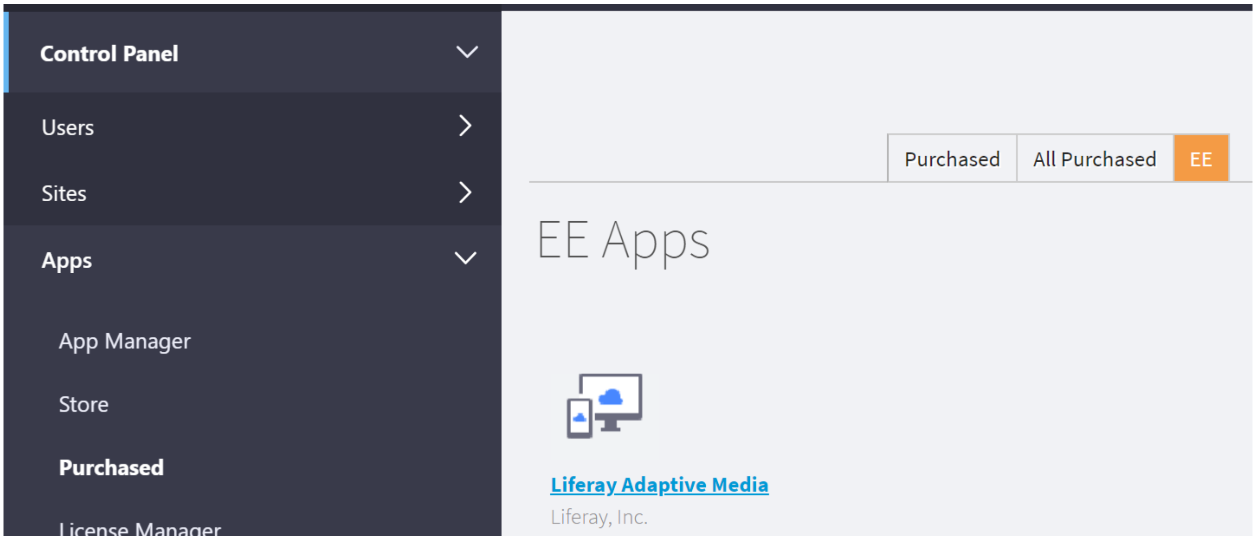

Figure 85.10: The trial plugins are available as plugins already purchased.

4. Once signed into the Store, click on the *Purchased* link, and then click on the *EE* subtab.

   Here you can see a list of Liferay DXP plugins that are installed, as well as options to update or install certain plugins. To install a plugin, click the *Install* button located to the far right of the respective plugin. Similarly, click the *Update* button to update a plugin.

Note: At a later date, the Marketplace website on `liferay.com` will be accessible to Liferay DXP trial customers. For now, please access subscription apps through the portal installation itself.

FAQ

Q: Where are the *Liferay DXP Trial Plugins*?

A: There is no such thing. The Liferay DXP plugins in Liferay Marketplace are the same ones that you get to try out with your Liferay DXP trial license for your portal. The Liferay DXP license (trial or official Liferay DXP subscriber) gives you access to the Liferay DXP plugins. Also, there is no difference code-wise or release-wise between a Liferay DXP trial installation and a regular Liferay DXP non-trial installation. The only difference is the license.

Q: Why can't I go to liferay.com/marketplace? Why can't I *purchase* from the Marketplace site?

A: DXP trial users must use the Marketplace from within the product's Control Panel (instructions above). You do not need to *purchase* any DXP plugins because if you access Marketplace from within the Control Panel, Marketplace sees that you have a DXP license installed and gives access to DXP plugins. Official DXP subscription customers (i.e., non trial) can log into `liferay.com` with their designated DXP subscriber login and access all DXP plugins through the Marketplace website.

Q: Why are the plugins under the Purchased tab? If I click on the *DXP Marketplace* link, it does not let me get the DXP plugins.

A: Once you're signed into the Store, click on the *Purchased* tab, then click on the *EE* subtab.

Q: What happens when DXP trial customers become official Liferay Digital Experience subscribers?

A: They can still complete the above process, or they can also visit the Liferay Marketplace website.

Q: Do DXP trial customers get the DXP source code?

A: No, they can only install the plugin. The DXP source code becomes available once they are official Liferay DXP Enterprise subscribers.

Q: Can this process of installing DXP plugins be used from Liferay Portal CE (Community Edition)?

A: No, the Marketplace must detect that you are running Liferay DXP.

85.12 Activating Liferay DXP

There are two ways to activate your Liferay DXP instance:

- With an XML activation key that you request and receive from Liferay Support.
- Online activation through Liferay Connected Services (LCS). Liferay DXP 7.0 introduced LCS as a way to activate Liferay DXP instances. LCS can also install fix packs, monitor each instance's performance, and help administrators automatically manage Liferay DXP subscriptions. See the LCS documentation for instructions on activating your instances with LCS.

Note: You must use LCS for activation of Elastic subscriptions. Otherwise, you don't have to use LCS for activation. You can instead request an XML activation key from Liferay Support.

85.13 Setting Up Marketplace

Liferay Marketplace is more than just a store for Liferay applications. Under the hood, it provides both the store and Liferay DXP's application deployment features. For this reason, you must ensure that Marketplace can run and configure itself.

Here are some scenarios to work around to ensure Marketplace works successfully:

- Server is Firewalled without Access to the Internet
- Limited Database Access

The firewall scenario is discussed first.

Server is Firewalled without Access to the Internet

Your server might be behind a firewall that prevents access to the Internet. Or your security policy might not allow direct download and installation from the Internet. In these cases, you have two options:

1. From an Internet-enabled computer, download the Marketplace plugin. Then allow Liferay DXP to auto deploy it by dropping the downloaded `.lpkg` file into the `deploy` folder in Liferay Home.

2. Alternately, once you have the downloaded `.lpkg` file, deploy it using the App Manager.

Next you'll learn how to work around database access restrictions.

Limited Database Access

Some production environments do not have the necessary database permissions for Liferay DXP, apps, modules, and plugins to maintain their tables. In these cases:

1. Grant the Liferay DXP database user temporary full rights to the database.
2. Install Liferay DXP and start it so that it populates its database.
3. Once the database is created, remove the permissions for creating tables and dropping tables from the Liferay DXP database user.

See the database server and new database instructions for more information. Note that many sophisticated Liferay DXP apps—not just the Marketplace app—require new tables when deployed. If your environment restricts database access, you may need to repeat the above steps whenever you deploy a new app.

You've prepared Liferay DXP for installing Marketplace and additional apps.

Chapter 86

Installing a Search Engine

A search engine is a critical component of your Liferay DXP installation. If you're here, you probably know the basics already and want to configure a search engine for your Liferay DXP deployment.

Liferay DXP ships with Elasticsearch, a highly scalable, full-text search engine. Elasticsearch is well-supported and almost certainly meets any search and indexing need you have. For deployment settings, learn to configure a standalone or remote Elasticsearch server or cluster here.

Solr is also supported in Liferay DXP.

86.1 Choosing a Search Engine

Elasticsearch and Solr are both supported, but there are some differences in how they work with Liferay DXP. In certain cases, you must choose Elasticsearch.

If you answer *yes* to either of these questions, you must choose Elasticsearch:

1. You're using Liferay Commerce, Liferay's commerce solution.

2. Your custom search code requires the use of the `TermsSetFilter` API or the Geolocation APIs that are implemented in the Liferay Connector to Elasticsearch.

@commerce@ requires the `TermsSetFilter` implementation available in the Elasticsearch connector, so you must use Elasticsearch if you're using @commerce@.

Both of these Elasticsearch-only developer features are not currently implemented in the Liferay Connector to Solr, but may be added in the future. If you must use either of those features in your search solution's code, use Elasticsearch. If you're using Liferay Commerce, use Elasticsearch. Otherwise, feel free to use Elasticsearch or Solr to index your portal content.

Another factor to consider in your search engine selection is JDK version. The search engine and Liferay DXP must use the same JDK version and distribution (e.g., Oracle Open JDK 1.8.0_201). Consult the Elasticsearch compatibility matrix and the Liferay DXP compatibility matrix to learn more about supported JDK distributions and versions. This consideration is not necessary for Solr, because no JVM level serialization happens between the servers. All communication occurs at the HTTP level.

Chapter 87

Elasticsearch

Elasticsearch is an open source, highly scalable, full-text search and analytics engine.

By default, Elasticsearch runs as an embedded search engine, which is useful for development and testing but is not supported in production. In production environments you must run Elasticsearch in remote mode, as a separate server or cluster. This guide walks you through the process of configuring Elasticsearch in remote mode.

If you'd rather use Solr, it's also supported. See the documentation on Installing Solr if you're interested.

To get up and running quickly with Elasticsearch as a remote server, refer to the Installing Elasticsearch article. Those are basic instructions for installing and configuring Elasticsearch in a single server environment. This article includes more details and information on clustering and tuning Elasticsearch. Here, you'll learn to configure your existing Elasticsearch installation for use in production environments.

If you've come here looking for information on search engines in general, or the low level search infrastructure of Liferay DXP, refer instead to the developer tutorial Introduction to Liferay Search (not yet written).

These terms are useful to understand as you read this guide:

- *Elasticsearch Home* refers to the root folder of your unzipped Elasticsearch installation (for example, `elasticsearch-6.5.1`).
- *Liferay Home* refers to the root folder of your Liferay DXP installation. It contains the `osgi`, `deploy`, `data`, and `license` folders, among others.

87.1 Embedded vs. Remote Operation Mode

When you start Liferay DXP, this message is displayed in the log:

```
2018-12-10 16:20:32.987 WARN  [Elasticsearch initialization thread][EmbeddedElasticsearchConnection:288] Liferay is configured to use embedded Elasti
```

When you install Liferay DXP, Elasticsearch is already embedded. In embedded mode, Elasticsearch search runs in the same JVM to make it easy to test-drive with minimal configuration. Running both servers in the same process has drawbacks:

- Elasticsearch must use the same JVM options as Liferay DXP.
- Liferay DXP and Elasticsearch compete for resources.

Note: While it's not a supported production configuration, installing Kibana to monitor the embedded Elasticsearch server is useful during development and testing. Just be aware that you must install the OSS only Kibana build.

You wouldn't run an embedded database like HSQL in production, and you shouldn't run Elasticsearch in embedded mode in production either. Instead, run Elasticsearch in *remote operation mode*, as a standalone server or cluster of server nodes.

87.2 Troubleshooting Elasticsearch

Sometimes things don't go as planned. If you've set up Liferay DXP with Elasticsearch in remote mode, but Liferay DXP can't connect to Elasticsearch, check these things:

Cluster name: The value of the `cluster.name` property in Elasticsearch must match the `clusterName` property you configured for Liferay's Elasticsearch adapter.

Transport address: The value of the `transportAddress` property in the Elasticsearch adapter must match the port where Elasticsearch is running. If Liferay DXP is running in embedded mode, and you start a standalone Elasticsearch node or cluster, it detects that port `9300` is taken and switches to port `9301`. If you then set Liferay's Elasticsearch adapter to remote mode, it continues to look for Elasticsearch at the default port (`9300`).

The following articles cover the Liferay Connector to Elasticsearch's configuration options in more detail.

87.3 Preparing to Install Elasticsearch

By default, 7.0 and its embedded Elasticsearch engine run in the same JVM. Although this enables out-of-the-box search, it's only supported for development. For production use, Elasticsearch must run in a separate JVM. See the installation guide for information on installing a remote Elasticsearch cluster.

Because search engines benefit heavily from caching, their JVM memory profiles differ substantially from those of a JVM running Liferay DXP. Therefore, the two applications should always be kept separate in production environments.

The following sections provide a synopsis of Elasticsearch configurations for 7.0. Prior to deployment, we strongly recommend reading Elastic's documentation on production deployment.

Sizing Your Deployment

When sizing your Elasticsearch deployment, carefully consider CPU, memory, disk, and network capacity. To scale effectively and avoid using lots of machines, deploy Elasticsearch on medium to large machines (for example, machines with two to eight CPUs). Avoid running multiple Elasticsearch JVMs on the same operating system.

CPU

We recommend allocating at least eight total CPU cores to the Elasticsearch engine, assuming only one Elasticsearch JVM is running on the machine.

Memory

At least 16 GB of memory is recommended, with 64 GB preferred. The precise memory allocation required depends on how much data is indexed. For index sizes 500 GB to 1 TB, 64 GB of memory suffices.

Disk

Search engines store their indexes on disk, so disk I/O capacity can impact search performance. Deploy Elasticsearch on SSD whenever possible. Otherwise use high-performance traditional hard disks (for example, 15k RPM). In either case, consider using RAID 0.

Avoid using Network Attached Storage (NAS) whenever possible as the network overhead can be large. If you're using public cloud infrastructure like Amazon Web Services, use instance local storage instead of network storage, such as Elastic Block Store (EBS).

Maintain 25 percent more disk capacity than the total size of your indexes. If your index is 60 GB, make sure you have at least 75 GB of disk space available. To estimate the disk space you need, you can index a representative sample of your production content and multiply that size by the fraction of your production content that it represents. For example, index 25 percent of your production content and then multiply the resulting index size by four. Keep in mind that indexing a 1 MB file doesn't result in 1 MB of disk space in the search index.

Cluster Size

While Liferay DXP can work with an Elasticsearch cluster comprised of one or two nodes, the minimum cluster size recommended by Elastic for fault tolerance is three nodes.

Networking

Elasticsearch relies on clustering and sharding to deliver fast, accurate search results, and thus requires a fast and reliable network. Most modern data centers provide 1 GbE or 10 GbE between machines.

Elasticsearch doesn't support multi-data center deployments.

87.4 Configuring the Liferay Elasticsearch Connector

For detailed Elasticsearch configuration information, refer to the Elasticsearch documentation.

The name of your Elasticsearch cluster is important. When you're running Elasticsearch in remote mode, the cluster name is used by Liferay DXP to recognize the Elasticsearch cluster. To learn about setting the Elasticsearch cluster name on the Liferay DXP side, refer below to the section called Configuring the Liferay Elasticsearch Connector.

Note: The `http.enabled` setting in Elasticsearch corresponds to the `httpEnabled` setting in the Liferay Connector to Elasticsearch 6 application. As this setting was deprecated in Elasticsearch

6.3, the connector's corresponding setting is now also deprecated. This setting was only used for configuring the embedded Elasticsearch server, so its deprecation should have minimal impact to production deployments.

Elasticsearch's configuration files are written in YAML and kept in the `[Elasticsearch Home]/config` folder. The main configuration file is `elasticsearch.yml`, used for configuring Elasticsearch modules.

To set the name of the Elasticsearch cluster, open `[Elasticsearch Home]/config/elasticsearch.yml` and specify

```
cluster.name: LiferayElasticsearchCluster
```

Since `LiferayElasticsearchCluster` is the default name given to the cluster, this would work just fine. Of course, you can name your cluster whatever you want (we humbly submit the recommendation `clustery_mcclusterface`).1 You can configure your node name using the same syntax (setting the `node.name` property).

If you'd rather work from the command line than in the configuration file, navigate to Elasticsearch Home and enter

```
./bin/elasticsearch --cluster.name clustery_mcclusterface --node.name nody_mcnodeface
```

Feel free to change the node name or the cluster name. Once you configure Elasticsearch to your liking, start it up.

Starting Elasticsearch

Start Elasticsearch by navigating to Elasticsearch Home and typing

```
./bin/elasticsearch
```

if you run Linux, or

```
\bin\elasticsearch.bat
```

if you run Windows.

To run as a daemon in the background, add the `-d` switch to either command:

```
./bin/elasticsearch -d
```

Once both Elasticsearch and Liferay DXP are installed and running, introduce them to each other.

Configuring the Liferay Elasticsearch Connector

The Elasticsearch connector provides integration between Elasticsearch and the portal. Before you configure the connector, make sure Elasticsearch is running.

There are two ways to configure the adapter:

1. Use the System Settings application in the Control Panel.

2. Manually create an OSGi configuration file.

It's convenient to configure the Elasticsearch adapter from System Settings, but this is often only possible during development and testing. If you're not familiar with System Settings, you can read about it here. Remember that you can generate configuration files for deployment to other systems by configuring System Settings, and then exporting the `.config` file with your configuration.

Configuring the Adapter in the Control Panel

To configure the Elasticsearch adapter from the System Settings application,

1. Start Liferay DXP.
2. Navigate to *Control Panel* → *Configuration* → *System Settings* → *Foundation*.
3. Find the *Elasticsearch* entry (scroll down and browse to it or use the search box) and click the Actions icon (⋮), then *Edit*.

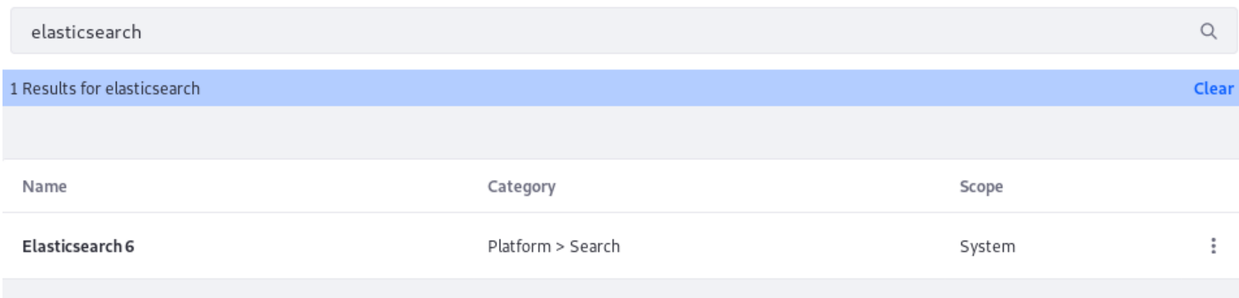

Figure 87.1: Use the System Settings application in Liferay DXP's Control Panel to configure the Elasticsearch adapter.

4. Make any edits to the configuration and click *Save*.

Note: If you switch operation modes (`EMBEDDED` → `REMOTE`), you must trigger a re-index. Navigate to *Control Panel* → *Configuration* → *Search*, and click *Execute* next to *Reindex all search indexes*.

Configuring the Adapter with an OSGi `.config` File

When preparing a system for production deployment, you want to use a repeatable deployment process. Therefore, it's best to use the OSGi configuration file, where your configuration is maintained in a controlled source.

Follow these steps to configure the Elasticsearch adapter using a configuration file:

1. Create the following file:

```
[Liferay_Home]/osgi/configs/com.liferay.portal.search.elasticsearch6.configuration.ElasticsearchConfiguration.config
```

2. Add configurations to the file, in the format `propertyName="Value"`. For example,

```
operationMode="REMOTE"
# If running Elasticsearch from a different computer:
#transportAddresses="ip.of.elasticsearch.node:9300"
# Highly recommended for all non-prodcution usage (e.g., practice, tests, diagnostics):
#logExceptionsOnly="false"
```

3. Start Liferay DXP or re-index if already running.

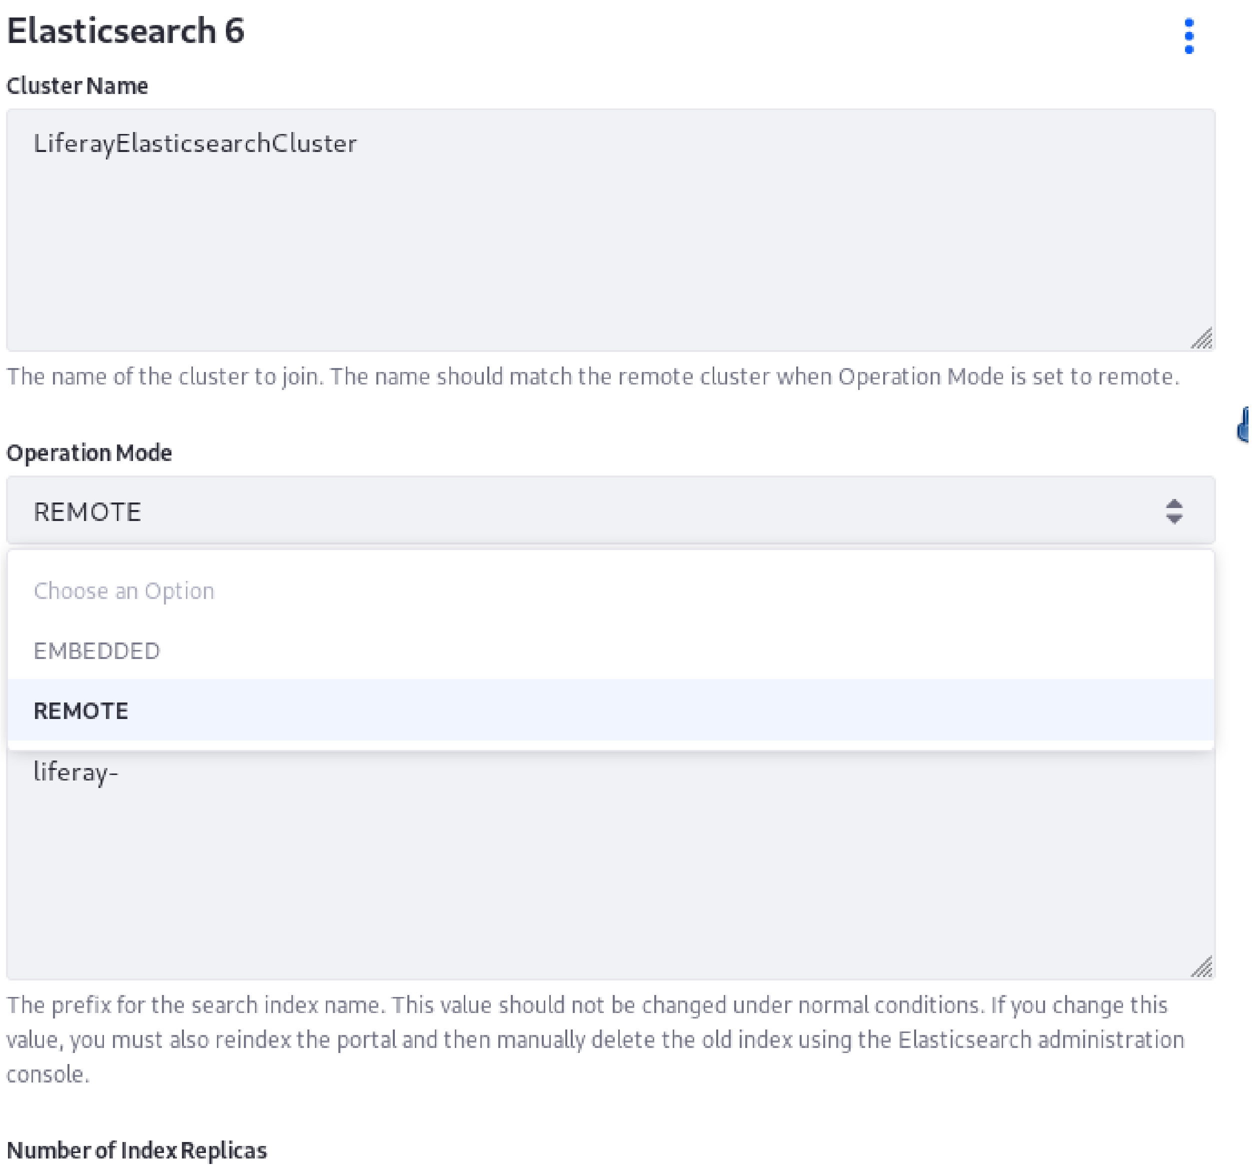

Figure 87.2: Set configurations for the Elasticsearch connector, like settings the Operation Mode to *Remote*.

As you can see from the System Settings entry for Elasticsearch, there are a lot more configuration options available that help you tune your system for optimal performance.

What follows here are some known good configurations for clustering Elasticsearch. These, however, can't replace the manual process of tuning, testing under load, and tuning again, so we encourage you to examine the Elasticsearch documentation and go through that process once you have a working configuration.

Configuring a Remote Elasticsearch Host

In production systems Elasticsearch and Liferay DXP are installed on different servers. To make Liferay DXP aware of the Elasticsearch cluster, set

```
transportAddresses=[IP address of Elasticsearch Node]:9300
```

Here's an example that sets the IP address of two nodes in the Elasticsearch cluster:

```
transportAddresses=["192.168.1.1:9300","192.168.1.2:9300"]
```

Set this in the Elasticsearch connector's OSGi configuration file. List as many or as few Elasticsearch nodes in this property as you want. This tells Liferay DXP the IP address or host name where search requests should be sent. If using System Settings, set the value in the *Transport Addresses* property.

Note: In an Elasticsearch cluster you can list the transport addresses for multiple Elasticsearch nodes as a comma-separated list in the `transportAddresses` property. If you set only one transport address, Liferay DXP loses contact with Elasticsearch if that node goes down.

On the Elasticsearch side, set the `network.host` property in your `elaticsearch.yml` file. This property simultaneously sets both the *bind host* (the host where Elasticsearch listens for requests) and the *publish host* (the host name or IP address Elasticsearch uses to communicate with other nodes). See here for more information.

Clustering Elasticsearch in Remote Operation Mode

Clustering Elasticsearch is easy. First, set `node.max_local_storage_nodes` to be something greater than `1`. When you run the Elasticsearch start script, a new local storage node is added to the cluster. If you want four nodes running locally, for example, run `./bin/elasticsearch` four times. See here for more information.

Configure the number of shards and replicas in the Elasticsearch 6 adapter, using the `indexNumberOfShards` and `indexNumberOfReplicas` properties to specify the number of primary shards and number of replica shards, respectively. Elasticsearch's default configuration works for a cluster of up to ten nodes, since the default number of shards is `5` and the default number of replica shards is `1`.

Note: Elasticsearch uses the Zen Discovery Module by default, which provides unicast discovery. Additionally, nodes in the cluster communicate using the Transport Module, through TCP. See the Elasticsearch documentation for the available properties (to be set in the `elasticsearch.yml` file), and the Liferay DXP Elasticsearch Adapter's settings for the adapter's available settings.

At a minimum, provide the list of hosts (as `host:port`) to act as gossip routers during unicast discovery in the `elasticsearch.yml`:

```
discovery.zen.ping.unicast.hosts: ["node1.ip.address", "node2.ip.address"]
```

For example,

```
discovery.zen.ping.unicast.hosts: ["10.10.10.5", "10.10.10,.5:9305"]
```

For more information on configuring an Elasticsearch cluster, see the documentation on Elasticsearch Index Settings.

Elasticsearch Connector System Settings, By Operation Mode

Some of the settings available for the Elasticsearch connector are applicable for only one operation mode (REMOTE or EMBEDDED). Refer to the table below:

Adapter Setting/Operation Mode | EMBEDDED | REMOTE | `clusterName` | x | x `operationMode` | x | x `indexNamePrefix` | x | x `indexNumberOfReplicas*` | x | x `indexNumberOfShards*` | x | x `bootstrapMlockAll` | x | - `logExceptionsOnly` | x | x `retryOnConflict` | x | x `discoveryZenPingUnicastHostsPort` | x | - `networkHost` | x | - `networkBindHost` | x | - `networkPublishHost` | x | - `transportTcpPort` | x | - `transportAddresses` | - | x `clientTransportSniff` | - | x `clientTransportIgnoreClusterName` | - | x `clientTransportPingTimeout*` | - | x `clientTransportNodesSamplerInterval` | - | x `httpEnabled` | x | - `httpCORSEnabled` | x | - `httpCORSAllowOrigin` | x | - `httpCORSConfigurations` | x | - `additionalConfigurations` | x | x `additionalIndexConfigurations` | x | x `additionalTypeMappings` | x | x `overrideTypeMappings` | x | x `syncSearch` | x | -

* **Note:** Available in the Connector to Elasticsearch 6 only.

1 This is, of course, a nod to all those fans of Boaty Mcboatface.

87.5 Advanced Configuration of the Liferay Elasticsearch Connector

The default configurations for Liferay's Elasticsearch adapter module are set in a Java class called `ElasticsearchConfiguration`.

While the Elasticsearch adapter has a lot of configuration options out of the box, you might find an Elasticsearch configuration you need that isn't provided by default. In this case, add the configuration options you need. If something is configurable for Elasticsearch, it's configurable using the Elasticsearch adapter.

Adding Settings and Mappings to the Liferay Elasticsearch Adapter

Think of the available configuration options as being divided into two groups: the most common ones that are easily configured, and more complex configurations requiring a more brute-force approach. If a necessary setting isn't available by default, you can still configure it with the Liferay Elasticsearch adapter. You'll just need to use one or more of the `additionalConfigurations`, `additionalIndexConfigurations`, or `additionalTypeMappings`, and `overrideTypeMappings` settings.

Additional Configurations

The `additionalConfigurations` configuration defines extra settings (in YAML) for the embedded Elasticsearch. This is only useful for testing environments using the embedded Elasticsearch server. Any node settings normally set in `elasticsearch.yml` can be declared here. See the Elasticsearch documentation for a description of all possible node settings.

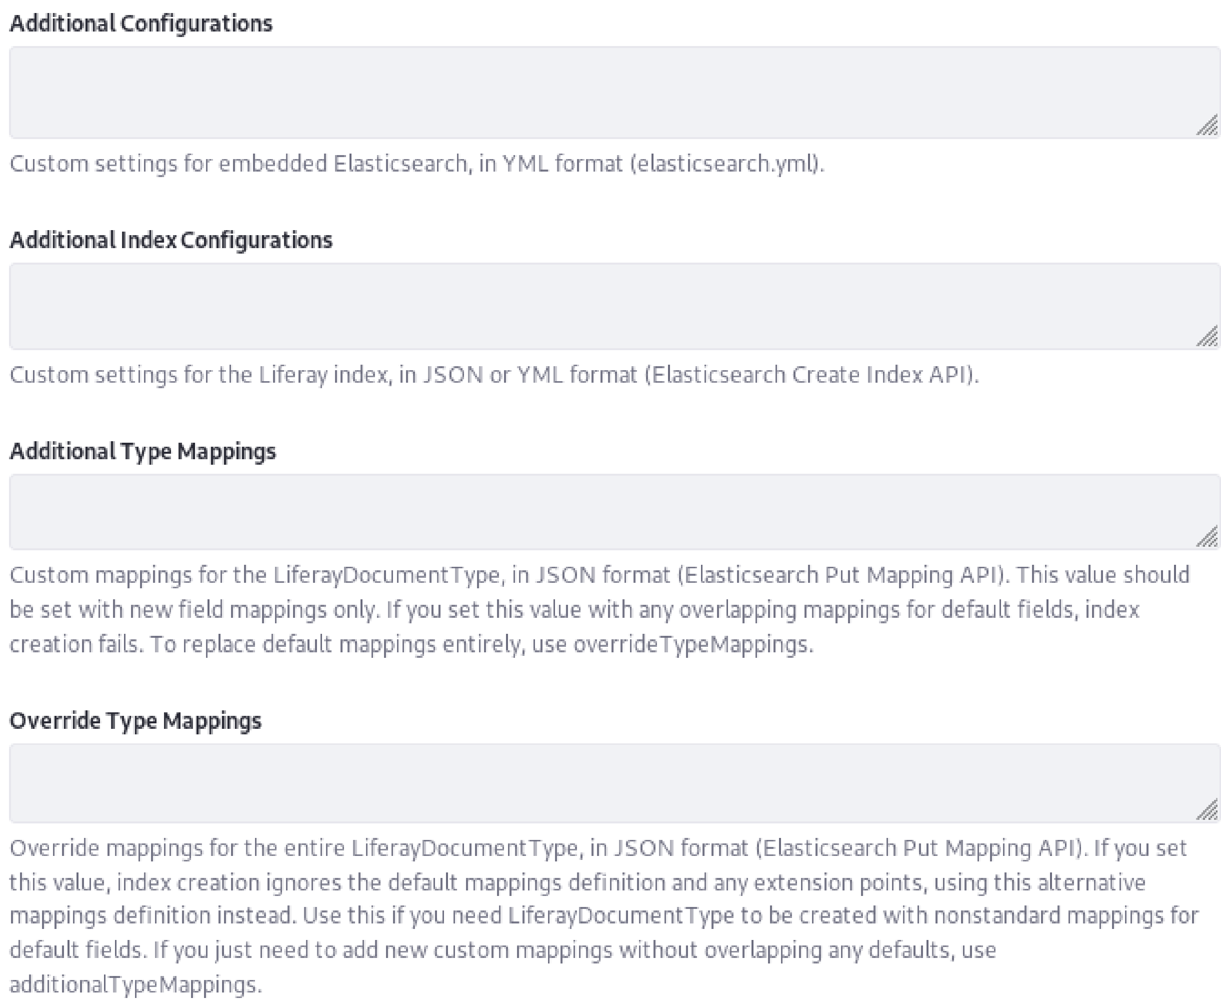

Figure 87.3: You can add Elasticsearch configurations to the ones currently available in System Settings.

Adding Index Configurations

The `additionalIndexConfigurations` configuration defines extra settings (in JSON or YAML) that are applied to the Liferay DXP index when it's created. For example, you can create custom analyzers and filters using this setting. For a complete list of available settings, see the Elasticsearch reference.

Here's an example that shows how to configure analysis that can be applied to a field or a dynamic template (see below for an example application to a dynamic template).

```
{
    "analysis": {
        "analyzer": {
            "kuromoji_liferay_custom": {
                "filter": [
                    "cjk_width",
                    "kuromoji_baseform",
                    "pos_filter"
                ],
                "tokenizer": "kuromoji_tokenizer"
            }
        },
```

```
        "filter": {
            "pos_filter": {
                "type": "kuromoji_part_of_speech"
            }
        }
    }
}
```

Adding Type Mappings

`additionalTypeMappings` defines extra mappings for the `LiferayDocumentType` type definition. These are applied when the index is created. Add the mappings in using JSON syntax. For more information see here and here. Use `additionalTypeMappings` for new field (`properties`) mappings and new dynamic templates, but don't try to override existing mappings. If any of the mappings set here overlap with existing mappings, index creation fails. Use `overrideTypeMappings` to replace default mappings.

As with dynamic templates, you can add sub-field mappings to Liferay DXP's type mapping. These are referred to as properties in Elasticsearch.

```
{
    "LiferayDocumentType": {
        "properties": {
            "fooName": {
                "index": "true",
                "store": "true",
                "type": "keyword"
            }
        }
    }
}
```

See here for more details on Elasticsearch's field datatypes.

The above example shows how a `fooName` field might be added to Liferay DXP's type mapping. Because `fooName` is not an existing property in the mapping, it works fine. If you try to override an existing property mapping, index creation fails. Instead use the `overrideTypeMappings` setting to override `properties` in the mapping.

To see that your additional mappings have been added to the `LiferayDocumentType`, use `curl` to access this URL after saving your additions and re-indexing:

```
curl http://[HOST]:[ES_PORT]/liferay-[COMPANY_ID]/_mapping/LiferayDocumentType?pretty
```

Here's what it would look like for an Elasticsearch instance running on `localhost:9200`, with a Liferay DXP Company ID of `20116`:

```
curl http://localhost:9200/liferay-20116/_mapping/LiferayDocumentType?pretty
```

In the above URL, `liferay-20116`is the index name. Including it indicates that you want to see the mappings that were used to create the index with that name.

Overriding Type Mappings

Use `overrideTypeMappings` to override Liferay DXP's default type mappings. This is an advanced feature that should be used only if strictly necessary. If you set this value, the default mappings used to define the Liferay Document Type in Liferay DXP source code (for example, `liferay-type-mappings.json`) are ignored entirely, so include the whole mappings definition in this property, not just the segment you're modifying. To make a modification, find the entire list of the current mappings being used to create the index by navigating to the URL

```
http://[HOST]:[ES_PORT]/liferay-[COMPANY_ID]/_mapping/LiferayDocumentType?pretty
```

Copy the contents in as the value of this property (either into System Settings or your OSGi configuration file). Leave the opening curly brace {, but delete lines 2-4 entirely:

```
"liferay-[COMPANY_ID]": {
    "mappings" : {
        "LiferayDocumentType" : {
```

Then, from the end of the mappings, delete the concluding three curly braces.

```
        }
    }
}
```

Now modify whatever mappings you'd like. The changes take effect once you save the changes and trigger a re-index from Server Administration.

Here's a partial example, showing a dynamic template that uses the analysis configuration from `additionalIndexConfigurations` to analyze all string fields that end with `_ja`. You'd include this with all the other default mappings, replacing the provided `template_ja` with this custom one:

```
{
    "LiferayDocumentType": {
        "dynamic_templates": [
            {
                "template_ja": {
                    "mapping": {
                        "analyzer": "kuromoji_liferay_custom",
                        "index": "analyzed",
                        "store": "true",
                        "term_vector": "with_positions_offsets",
                        "type": "string"
                    },
                    "match": "\\w+_ja\\b|\\w+_ja_[A-Z]{2}\\b",
                    "match_mapping_type": "string",
                    "match_pattern": "regex"
                }
                ...
            }
        ]
    }
}
```

Multi-line YAML Configurations

If you configure the settings from the last section using an OSGi configuration file, you might find yourself needing to write YAML snippets that span multiple lines. The syntax for that is straightforward and just requires appending each line with `\n\`, like this:

```
additionalConfigurations=\
                    cluster.routing.allocation.disk.threshold_enabled: false\n\
                    cluster.service.slow_task_logging_threshold: 600s\n\
                    index.indexing.slowlog.threshold.index.warn: 600s\n\
                    index.search.slowlog.threshold.fetch.warn: 600s\n\
                    index.search.slowlog.threshold.query.warn: 600s\n\
                    monitor.jvm.gc.old.warn: 600s\n\
                    monitor.jvm.gc.young.warn: 600s
```

From simple configurations to overriding existing type mappings, Elasticsearch and Liferay's connector to Elasticsearch are configurable.

87.6 Elasticsearch Connector Settings: Reference

Elasticsearch is the default search engine for 7.0. The *Liferay Foundation* suite includes an adapter for Elasticsearch called *Liferay Connector to Elasticsearch 6*. The adapter is configurable through System Settings or an OSGi configuration file named `com.liferay.portal.search.elasticsearch6.configuration.Elasticsearc` and deployed to `[Liferay_Home]/osgi/configs`.

The list below is all the configuration settings for Liferay's default Elasticsearch adapter, in the order they appear in the System Settings application (The *Elasticsearch 6* entry under the *Search* category):

`clusterName=LiferayElasticsearchCluster` A String value that sets the name of the cluster to integrate with. This name should match the remote cluster when Operation Mode is set to remote. (See also: remote operation mode)

`operationMode=EMBEDDED` There are two operation modes you can choose from: EMBEDDED or REMOTE. Set to REMOTE to connect to a remote standalone Elasticsearch cluster. Set to EMBEDDED to start Liferay with an internal Elasticsearch instance. Embedded operation mode is unsupported for production environments.

`indexNamePrefix=liferay-` Set a String value to use as the prefix for the search index name. The default value should not be changed under normal conditions. If you change it, you must also perform a *reindex all* operation for the portal and then manually delete the old index using the Elasticsearch administration console.

`indexNumberOfReplicas=` Set the number of replicas for each index. If left unset, no replicas are used. A full reindex is required to make changes take effect.

`indexNumberOfShards=` Set the number of index shards to use when a Liferay index is created. If left unset, a single shard is used. A full reindex is required to make changes take effect.

`bootstrapMlockAll=false` A boolean setting that, when set to `true`, tries to lock the process address space into RAM, preventing any Elasticsearch memory from being swapped out (see here) for more information)

`logExceptionsOnly=true` A boolean setting that, when set to true, only logs exceptions from Elasticsearch, and does not rethrow them.

`retryOnConflict=5` Set an int value for the number of retries to attempt if a version conflict occurs because the document was updated between getting it and updating it (see here for more information).

`discoveryZenPingUnicastHostsPort=9300-9400` Set a String value for the range of ports to use when building the value for discovery.zen.ping.unicast.hosts. Multiple Elasticsearch nodes on a range of ports can act as gossip routers at the same computer (see here for more information).

`networkHost=` Set this String value to instruct the node to bind to this hostname or IP address and publish (advertise) this host to other nodes in the cluster. This is a shortcut which sets the bind host and the publish host at the same time (see here for more information).

`networkBindHost=` Set the String value of the network interface(s) a node should bind to in order to listen for incoming requests (see here for more information).

`networkPublishHost=` Set the String value of a single interface that the node advertises to other nodes in the cluster, so that those nodes can connect to it (see here for more information).

`transportTcpPort=` Set the String value for the port to bind for communication between nodes. Accepts a single value or a range (see here for more information).

`transportAddresses=localhost:9300` Set the String values for the addresses of the remote Elasticsearch nodes to connect to. This value is required when Operation Mode is set to remote (see here for more information). Specify as many or few nodes as you see fit.

`clientTransportSniff=true` Set this booleant to true to enable cluster sniffing and dynamically discover available data nodes in the cluster (see here for more information).

`clientTransportIgnoreClusterName=false` Set this boolean to true to ignore cluster name validation of connected nodes (see here for more information).

`clientTransportPingTimeout=` The time (in seconds) the client node waits for a ping response from a node. If unset, the default Elasticsearch `client.transport.ping_timeout` is used.

`clientTransportNodesSamplerInterval=` Set this String value to instruct the client node on how often to sample / ping the nodes listed and connected (see here for more information).

`httpEnabled=true` Set this boolean to false to disable the http layer entirely on nodes which are not meant to serve REST requests directly. As this setting was deprecated in Elasticsearch 6.3, the connector's corresponding setting is now also deprecated. This setting was only used for configuring the embedded Elasticsearch server, so its deprecation should have minimal impact to production deployments.

`httpCORSEnabled=true` Set this boolean to false to disable cross-origin resource sharing, i.e. whether a browser on another origin can do requests to Elasticsearch. If disabled, web front end tools like elasticsearch-head may be unable to connect (see here for more information).

`httpCORSAllowOrigin=/https?:\\/\\/localhost(:[0-9]+)?/` Set the String origins to allow when HTTP CORS is enabled (see here for more information).

`httpCORSConfigurations=` Set the String values for custom settings for HTTP CORS, in YML format (elasticsearch.yml) (see here for more information).

`additionalConfigurations=` Set the String values for custom settings for embedded Elasticsearch, in YML format. See: Adding Settings to the Liferay Elasticsearch Adapter

`additionalIndexConfigurations=` Set the String values for custom settings for the Liferay index, in JSON or YML format (refer to the Elasticsearch Create Index API for more information). See: Adding Settings to the Liferay Elasticsearch Adapter

`additionalTypeMappings=` Set the String values for custom mappings for the `LiferayDocumentType`, in JSON format (refer to the Elasticsearch Put Mapping API for more information) See: Adding Settings to the Liferay Elasticsearch Adapter

`overrideTypeMappings=` Settings here override Liferay DXP's default type mappings. This is an advanced feature that should be used only if strictly necessary. If you set this value, the default mappings used to define the Liferay Document Type in Liferay DXP source code (for example, `liferay-type-mappings.json`) are ignored entirely, so include the whole mappings definition in this property, not just the segment you're modifying.

`syncSearch=true` If enabled, search runs on the invoker thread rather than in Elasticsearch's search thread pool.

The following settings are only available in the Elasticsearch 6 adapter:

Configurations only Affecting the Embedded Elasticsearch Server

These settings (defined above) are only meant to use while configuring the embedded Elasticsearch server. Configuring these will elicit no effect on remote Elasticsearch installations:

- `bootstrapMlockAll`

- `discoveryZenPingUnicastHostsPort`
- `networkHost`
- `networkBindHost`
- `networkPublishHost`
- `transportTcpPort`
- `httpEnabled`
- `httpCORSEnabled`
- `httpCORSAllowOrigin`
- `httpCORSConfigurations`
- `syncSearch`

You can easily configure these settings in the System Setting application, or as mentioned above, you can specify them in a deployable OSGi `.config` file.

87.7 Tuning Elasticsearch

Since search engines benefit heavily from caching, their JVM memory profiles are substantially different from those of a JVM focused on serving content and web views (e.g., a JVM running Liferay DXP). In production environments, search engines and Liferay DXP should always be on separate JVMs.

The following sections provide a synopsis of Elasticsearch configurations. Prior to deployment, we strongly recommend reading Elastic's documentation on production deployment.

You'll learn how to configure these settings:

- JVM
- File System
- Scale

JVM

The JVM vendor and version must be the same for the Elasticsearch server and the Liferay DXP server. In general, you should allocate 45 percent of the available system memory to Elasticsearch, up to a maximum of 31 GB. Configure heap sizing by setting the `ES_HEAP_SIZE` environment variable.

File System

Configure your operating system for at least 64,000 file descriptors (the default Linux value is 1024). Since Elasticsearch uses NioFS and MMapFS, ensure there is sufficient virtual memory available for memory-mapped files. Consult your system administrator for information on how to configure these values.

Tuning and Scaling an Elasticsearch Cluster

Proper scaling and tuning of an Elasticsearch cluster primarily depends on the type of indexes it holds and how they're intended to be used. Since Liferay DXP is a flexible development platform, no two applications index and search for data in exactly the same way. Read the definitive Elasticsearch guide, and understand the differences between indexing-intensive applications and search-intensive applications. Then you'll be able to predict usage patterns for your Liferay DXP indexes and design the optimally scaled and tuned cluster.

Once you determine the appropriate number of shards and replicas, configure them in the Liferay Connector to Elasticsearch module, using these settings:

- `indexNumberOfReplicas` corresponds to Elasticsearch's `number_of_replicas` property.
- `indexNumberOfShards` corresponds to Elasticsearch's `number_of_shards` property.

Tune, scale, and prosper.

87.8 Backing Up Elasticsearch

Elasticsearch replicas protect against a node going down, but they won't help you with a catastrophic failure. Only good backup practices can help you then.

Back up and restore your Elasticsearch cluster in three steps:

1. Configure a repository
2. Make a snapshot of the cluster
3. Restore from the snapshot

For more detailed information, refer to the Elasticsearch administration guide, and in particular to the documentation on the Snapshot/Restore module.

Creating a Repository

First create a repository to store your snapshots. Several repository types are supported:

- Shared file system, such as a Network File System or NAS
- Amazon S3
- HDFS (Hadoop Distributed File System)
- Azure Cloud

If using a shared file system repository type, first register the path to the shared file system in each node's `elasticsearch.yml` using the path.repo setting.

```
path.repo: ["path/to/shared/file/system/"]
```

Once the path to the folder hosting the repository is registered (make sure the folder exists), create the repository with a PUT command:

```
curl -X PUT "localhost:9200/_snapshot/test_backup" -H 'Content-Type: application/json' -d'
{
  "type": "fs",
  "settings": {
    "location": "/path/to/shared/file/system/"
  }
}
'
```

Replace `localhost:9200` with the proper `hostname:port` combination for your system, replace `test_backup` with the name of the repository to create, and use the absolute path to your shared file system in the `location`.

If the repository is set up successfully, you should see this message:

```
{"acknowledged":true}
```

Once the repository exists, you can start creating snapshots.

Taking Snapshots of the Cluster

The easiest snapshot approach is to create a snapshot of all the indexes in your cluster. To snapshot everything, enter

```
curl -XPUT localhost:9200/_snapshot/test_backup/snapshot_1
```

If `{"accepted":true}` appears in the terminal, the snapshot was a success.

It's possible to be more selective when taking snapshots. For example, if you use LES Monitoring, you can exclude the monitoring indexes. Explicitly declare the indexes to include in the snapshot:

```
curl -XPUT localhost:9200/_snapshot/test_backup/snapshot_2
{ "indices": "liferay-0,liferay-20116" }
```

Note: For a list of all the Elasticsearch indexes, use this command:

```
curl -X GET "localhost:9200/_cat/indices?v"
```

This shows the index metrics:

```
health status index         uuid                   pri rep docs.count docs.deleted store.size pri.store.size
green  open   liferay-20099 obqiNE1_SDqfuz7rincrGQ   1   0        195            0    303.1kb        303.1kb
green  open   liferay-47206 3YEjtye1S9OVT0i0EZcXcw   1   0          7            0     69.7kb         69.7kb
green  open   liferay-0     shBWwpkXRxuAmGEaE475ug   1   0        147            1    390.9kb        390.9kb
```

It's important to note that Elasticsearch uses a *smart snapshots* approach. To understand what that means, consider a single index. The first snapshot includes a copy of the entire index, while subsequent snapshots only include the delta between the first, complete index snapshot and the current state of the index.

Eventually you'll end up with a lot of snapshots in your repository, and no matter how cleverly you name the snapshots, you may forget what some snapshots contain. For this purpose, the Elasticsearch API provides getting information about any snapshot. For example:

```
curl -XGET localhost:9200/_snapshot/test_backup/snapshot_1
```

returns

```
{"snapshots":[
    {"snapshot":"snapshot_1",
    "uuid":"WlSjvJwHRh-xlAny7zeW3w",
    "version_id":6.50399,
    "version":"6.5.1",
    "indices":["liferay-20099","liferay-0","liferay-47206"],
    "state":"SUCCESS",
    "start_time":"2018-08-15T21:40:17.261Z",
```

```
    "start_time_in_millis":1534369217261,
    "end_time":"2018-08-15T21:40:17.482Z",
    "end_time_in_millis":1534369217482,
    "duration_in_millis":221,
    "failures":[],
    "shards":{
        "total":3,
        "failed":0,
        "successful":3

        }
    }
]}
```

There's lots of useful information here, including which indexes were included in the snapshot. If you want to get rid of a snapshot, use the DELETE command.

```
curl -XDELETE localhost:9200/_snapshot/test_backup/snapshot_1
```

You might trigger creation of a snapshot and regret it (for example, you didn't want to include all the indexes in the snapshot). If your snapshots contain a lot of data, this can cost time and resources. To cancel the ongoing creation of a snapshot, use the same DELETE command. The snapshot process is terminated and the partial snapshot is deleted from the repository.

Restoring from a Snapshot

What good is a snapshot if you can't use it to restore your search indexes in case of catastrophic failure? Use the `_restore` API to restore all the snapshot's indexes:

```
curl -XPOST localhost:9200/_snapshot/test_backup/snapshot_1/_restore
```

Restore only specific indexes from a snapshot by passing in the `indices` option, and rename the indexes using the `rename_pattern` and `rename_replacement` options:

```
curl -XPOST
localhost:9200/_snapshot/test_backup/snapshot_1/_restore
{
    "indices": "liferay-20116",
    "rename_pattern": "liferayindex_(.+)",
    "rename_replacement": "restored_liferayindex_$1"
}
```

This restores only the index named `liferay-20116index_1` from the snapshot. The `rename...` settings specify that the beginning `liferayindex_` are replaced with `restored_liferayindex_`, so `liferay-20116index_1` becomes `restored_liferay-20116index_1`.

As with the process for taking snapshots, an errant restored index can be canceled with the DELETE command:

```
curl -XDELETE localhost:9200/restored_liferay-20116index_3
```

Nobody likes catastrophic failure on a production system, but Elasticsearch's API for taking snapshots and restoring indexes can help you rest easy knowing that your search cluster can be restored if disaster strikes. For more details and options, read Elastic's documentation on the Snapshot and Restore Module.

Chapter 88

Installing Elasticsearch 6.1

Note: Elasticsearch 6.1.x is a supported search engine for 7.0, but it's not the newest supported version. If you're installing Elasticsearch for the first time, considering following the instructions for installing Elasticsearch 6.5.x instead, located in the root of the deployment guide.

Liferay DXP uses Elasticsearch to index its content. By default, it's installed as an embedded service. It works, but it's not a supported configuration for a production server. Feel free to use it while testing or developing, but when you're ready to put your site in production, you must run Elasticsearch as a standalone process. This is better anyway, because it frees you to design your infrastructure the way you want it. If you've got hardware or a VM to spare, you can separate your search infrastructure from Liferay DXP and reap some performance gains by putting search on a separate box. If you're more budget-conscious, you can still increase performance by running Elasticsearch in a separate, individually tunable JVM on the same box.

Installing Elasticsearch is pretty easy and takes only six steps:

1. Find the version of Elasticsearch that's embedded in the version of Liferay DXP you have, and then download that version from Elastic's website.
2. Install Elasticsearch by extracting its archive to the system where you want it to run.
3. Install some required Elasticsearch plugins.
4. Name your Elasticsearch cluster.
5. Configure Liferay DXP to connect to your Elasticsearch cluster.
6. Restart Liferay DXP and reindex your search indexes.

Note: Before continuing, make sure you have set the `JAVA_HOME` environment variable.

If you have multiple JDKs installed, make sure Elasticsearch and Liferay DXP are using the same version and distribution (e.g., Oracle Open JDK 1.8.0_201). You can specify this in `[Elasticsearch Home]/bin/elasticsearch.in.sh`:

```
JAVA_HOME=/path/to/java
```

Consult the Elasticsearch compatibility matrix and the Liferay DXP compatibility matrix to learn more about supported JDK distributions and versions.

Now you'll perform these steps, and when you're done, you'll have a production-ready instance of Liferay DXP up and running.

Step One: Find the Right Version of Elasticsearch

If Liferay DXP isn't running, start it.

Visit port 9200 on localhost to access the embedded Elasticsearch:

```
http://localhost:9200
```

A JSON document is returned that looks similar to this:

```
{
  "name" : "g0m223N",
  "cluster_name" : "LiferayElasticsearchCluster",
  "cluster_uuid" : "Ii6STs04Tg-XzTVV5h7M2Q",
  "version" : {
    "number" : "6.1.3",
    "build_hash" : "af51318",
    "build_date" : "2018-01-26T18:22:55.523Z",
    "build_snapshot" : false,
    "lucene_version" : "7.1.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "tagline" : "You Know, for Search"
}
```

The version of Elasticsearch that's running is the value of the `"number"` field. In this example, it's 6.1.3.

Now that you know the version of Elasticsearch you need, go to Elastic's website and download that version.

Step Two: Install Elasticsearch

Most of this step entails deciding where you want to run Elasticsearch. Do you want to run it on the same machine as Liferay DXP, or do you want to run it on its own hardware? The answer to this question comes down to a combination of the resources you have available and the size of your installation. Regardless of what you decide, either way you get the benefit of a separately tunable search infrastructure.

Once you have a copy of the right version of Elasticsearch, extract it to a folder on the machine where you want it running. That's it!.

Step Three: Install Elasticsearch Plugins

Install the following required Elasticsearch plugins:

- `analysis-icu`
- `analysis-kuromoji`
- `analysis-smartcn`
- `analysis-stempel`

To install these plugins, navigate to Elasticsearch Home and enter

```
./bin/elasticsearch-plugin install [plugin-name]
```

Replace *[plugin-name]* with the Elasticsearch plugin's name.

Step Four: Name Your Elasticsearch Cluster

A *cluster* in Elasticsearch is a collection of nodes (servers) identified as a cluster by a shared cluster name. The nodes work together to share data and workload. A one node cluster is discussed here; to create a multi-node cluster, please refer to Elastic's documentation.

Now that you've installed Elastic, it sits in a folder on your machine, which is referred to here as `[Elasticsearch Home]`. To name your cluster, you'll define the cluster name in both Elasticsearch and in Liferay DXP. First, define it in Elasticsearch. Edit the following file:

```
[Elasticsearch Home]/config/elasticsearch.yml
```

Uncomment the line that begins with `cluster.name`. Set the cluster name to whatever you want to name your cluster:

```
cluster.name: LiferayElasticsearchCluster
```

Of course, this isn't a very imaginative name; you may choose to name your cluster `finders_keepers` or something else you can remember more easily. Save the file.

Now you can start Elasticsearch. Run the executable for your operating system from the `[Elasticsearch Home]/bin` folder:

```
./elasticsearch
```

Elasticsearch starts, and one of its status messages includes a transport address:

```
[2018-04-03T15:34:19,784][INFO ][o.e.t.TransportService   ] [g0m223N] publish_address {127.0.0.1:9300}, bound_addresses {[::1]:9300}, {127.0.0.1:9300
```

Take note of this address; you'll need to give it to your Liferay DXP server so it can find Elasticsearch on the network.

Step Five: Configure Liferay DXP to Connect to your Elasticsearch Cluster

Now that you're ready to configure Liferay DXP, start it if you haven't already, log in, and then click on *Control Panel* → *Configuration* → *System Settings* → *Search*. Enter the term *elasticsearch* in the search bar and click the *Elasticsearch 6* entry from the list of settings. Now you can configure it. Here are the configuration options to change:

Cluster Name: Enter the name of the cluster as you defined it in Elasticsearch.

Operation Mode: Defaults to EMBEDDED. Change it to REMOTE to connect to a standalone Elasticsearch.

Transport Addresses: Enter a delimited list of transport addresses for Elasticsearch nodes. Here, you'll enter the transport address from the Elasticsearch server you started. The default value is `localhost:9300`, which will work.

When finished, click *Save*. You're almost done.

Step Six: Restart Liferay DXP and Reindex

Stop and restart Liferay DXP. When it's back up, log in as an administrative user and click on *Control Panel* → *Configuration* → *Search* and click the *Execute* button for *Reindex all search indexes*. When you do that, you should see some messages scroll up in the Elasticsearch log.

For more details refer to the Elasticsearch installation guide.

88.1 Configuring the Liferay Elasticsearch Connector (6.1)

For detailed Elasticsearch configuration information, refer to the Elasticsearch documentation.

The name of your Elasticsearch cluster is important. When you're running Elasticsearch in remote mode, the cluster name is used by Liferay DXP to recognize the Elasticsearch cluster. To learn about setting the Elasticsearch cluster name on the Liferay DXP side, refer below to the section called Configuring the Liferay Elasticsearch Connector.

Elasticsearch's configuration files are written in YAML and kept in the `[Elasticsearch Home]/config` folder:

- `elasticsearch.yml` is for configuring Elasticsearch modules
- `logging.yml` is for configuring Elasticsearch logging

To set the name of the Elasticsearch cluster, open `[Elasticsearch Home]/config/elasticsearch.yml` and specify

```
cluster.name: LiferayElasticsearchCluster
```

Since `LiferayElasticsearchCluster` is the default name given to the cluster, this would work just fine. Of course, you can name your cluster whatever you'd like (we humbly submit the recommendation `clustery_mcclusterface`).1 You can configure your node name using the same syntax (setting the `node.name` property).

If you'd rather work from the command line than in the configuration file, navigate to Elasticsearch Home and enter

```
./bin/elasticsearch --cluster.name clustery_mcclusterface --node.name nody_mcnodeface
```

Feel free to change the node name or the cluster name. Once you configure Elasticsearch to your liking, start it up.

Starting Elasticsearch

Start Elasticsearch by navigating to Elasticsearch Home and typing

```
./bin/elasticsearch
```

if you run Linux, or

```
\bin\elasticsearch.bat
```

if you run Windows.

To run as a daemon in the background, add the `-d` switch to either command:

```
./bin/elasticsearch -d
```

Once both Elasticsearch and Liferay DXP are installed and running, introduce Liferay DXP and Elasticsearch to each other.

Configuring the Liferay Elasticsearch Connector

The Elasticsearch connector provides integration between Elasticsearch and Liferay DXP. Before you configure the connector, make sure Elasticsearch is running.

There are two ways to configure the adapter:

1. Use the System Settings application in the Control Panel.

2. Manually create an OSGi configuration file.

It's convenient to configure the Elasticsearch adapter from System Settings, but this is often only possible during development and testing. If you're not familiar with System Settings, you can read about it here. Remember that you can generate configuration files for deployment to other systems by configuring System Settings, and then exporting the `.config` file with your configuration.

Configuring the Adapter in the Control Panel

Here's how to configure the Elasticsearch adapter from the System Settings application:

1. Start Liferay DXP.

2. Navigate to *Control Panel* → *Configuration* → *System Settings* → *Foundation*.

3. Find the *Elasticsearch* entry (scroll down and browse to it or use the search box) and click the Actions icon ( ), then *Edit*.

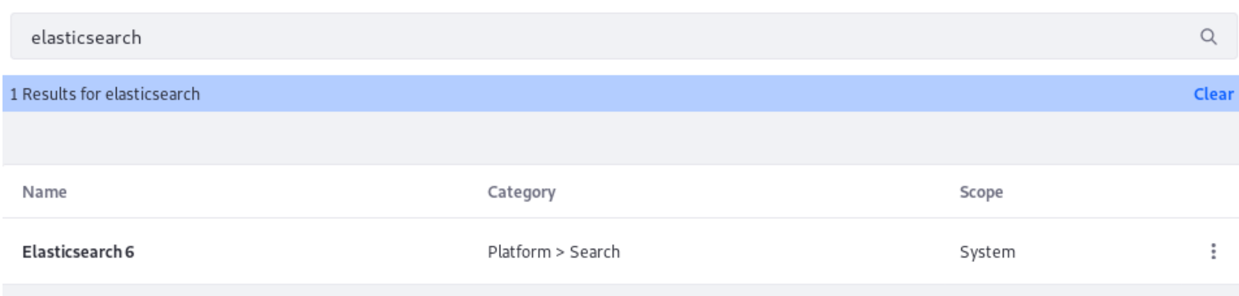

Figure 88.1: Use the System Settings application in Liferay DXP's Control Panel to configure the Elasticsearch adapter.

4. Make any edits to the configuration and click *Save*.

Note: If you switch operation modes (`EMBEDDED` → `REMOTE`), you must trigger a re-index. Navigate to *Control Panel* → *Configuration* → *Search*, and click *Execute* next to *Reindex all search indexes.*

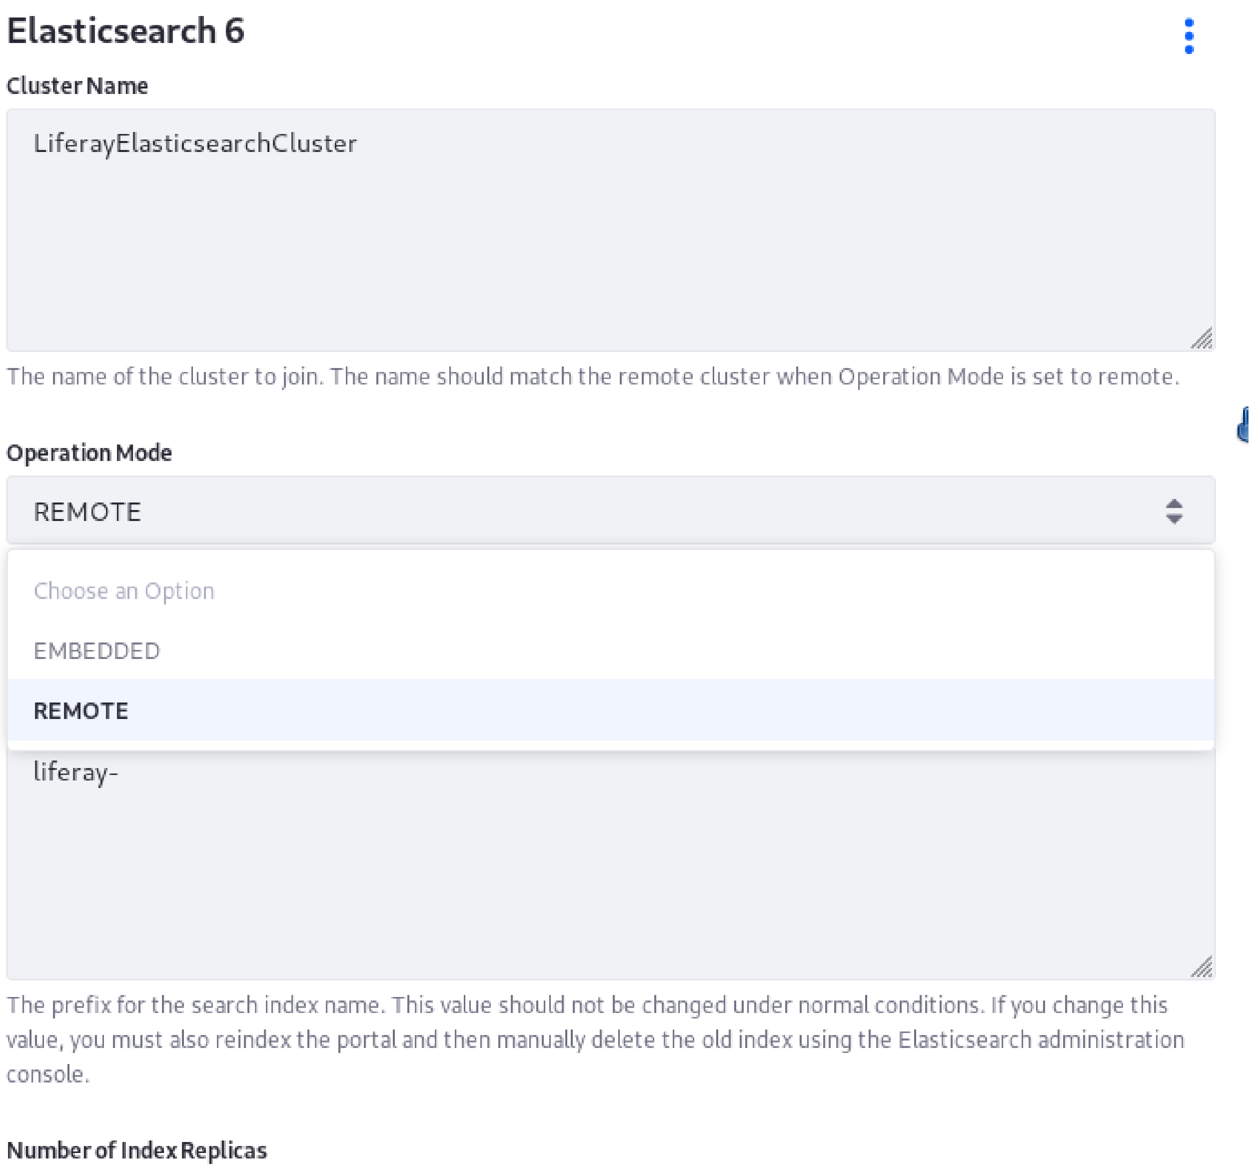

Figure 88.2: Set configurations for the Elasticsearch connector, like settings the Operation Mode to *Remote*.

Configuring the Adapter with an OSGi `.config` File

When preparing a system for production deployment, you should have a repeatable deployment process. Therefore, it's best to use the OSGi configuration file, where your configuration is maintained in a controlled source.

Follow these steps to configure the Elasticsearch adapter using an OSGi configuration file:

1. Create the following file:

```
[Liferay_Home]/osgi/configs/com.liferay.portal.search.elasticsearch6.configuration.ElasticsearchConfiguration.config
```

2. Add configurations to the file, in the format `propertyName="Value"`. For example,

```
operationMode="REMOTE"
# If running Elasticsearch from a different computer:
#transportAddresses="ip.of.elasticsearch.node:9300"
# Highly recommended for all non-prodcution usage (e.g., practice, tests, diagnostics):
#logExceptionsOnly="false"
```

3. Start Liferay DXP or re-index if already running.

As you can see from the System Settings entry for Elasticsearch, there are a lot more configuration options available that help you tune your system for optimal performance.

What follows here are some known good configurations for clustering Elasticsearch. These, however, can't replace the manual process of tuning, testing under load, and tuning again, so we encourage you to examine the Elasticsearch documentation and go through that process once you have a working configuration.

Configuring a Remote Elasticsearch Host

In production systems Elasticsearch and Liferay DXP are installed on different servers. To make Liferay DXP aware of the Elasticsearch cluster, set

```
transportAddresses=[IP address of Elasticsearch Node]:9300
```

Here's an example that sets the IP address of two nodes in the Elasticsearch cluster:

```
transportAddresses=["192.168.1.1:9300","192.168.1.2:9300"]
```

Set this in the Elasticsearch connector's OSGi configuration file. List as many Elasticsearch nodes in this property as you want. This tells Liferay DXP the IP address or host name where search requests should be sent. If using System Settings, set the value in the *Transport Addresses* property.

Note: In an Elasticsearch cluster you can list the transport addresses for multiple Elasticsearch nodes as a comma-separated list in the `transportAddresses` property. If you set only one transport address, Liferay DXP loses contact with Elasticsearch if that node goes down.

On the Elasticsearch side, set the `network.host` property in your `elaticsearch.yml` file. This property simultaneously sets both the *bind host* (the host where Elasticsearch listens for requests) and the *publish host* (the host name or IP address Elasticsearch uses to communicate with other nodes). See here for more information.

Clustering Elasticsearch in Remote Operation Mode

Clustering Elasticsearch is easy. First, set `node.max_local_storage_nodes` to be something greater than 1. When you run the Elasticsearch start script, a new local storage node is added to the cluster. If you want four nodes running locally, for example, run `./bin/elasticsearch` four times. See here for more information.

Configure the number of shards and replicas in the Elasticsearch 6 adapter, using the `indexNumberOfShards` and `indexNumberOfReplicas` properties to specify the number of primary shards and number of replica shards, respectively. Elasticsearch's default configuration works for a cluster of up to ten nodes, since the default number of shards is `5` and the default number of replica shards is `1`.

Note: Elasticsearch uses the Zen Discovery Module by default, which provides unicast discovery. Additionally, nodes in the cluster communicate using the Transport Module, through TCP. See the Elasticsearch documentation for the available properties (to be set in the `elasticsearch.yml` file), and the Liferay DXP Elasticsearch Adapter's settings for the adapter's available settings.

At a minimum, provide the list of hosts (as `host:port`) to act as gossip routers during unicast discovery in the `elasticsearch.yml`:

```
discovery.zen.ping.unicast.hosts: ["node1.ip.address", "node2.ip.address"]
```

For example,

```
discovery.zen.ping.unicast.hosts: ["10.10.10.5", "10.10.10,.5:9305"]
```

For more information on configuring an Elasticsearch cluster, see the documentation on Elasticsearch Index Settings.

Elasticsearch Connector System Settings, By Operation Mode

Some of the settings available for the Elasticsearch connector are applicable for only one operation mode (REMOTe or EMBEDDED). Refer to the table below:

Adapter Setting/Operation Mode | EMBEDDED | REMOTE | `clusterName` | x | x `operationMode` | x | x `indexNamePrefix` | x | x `indexNumberOfReplicas*` | x | x `indexNumberOfShards*` | x | x `bootstrapMlockAll` | x | - `logExceptionsOnly` | x | x `retryOnConflict` | x | x `discoveryZenPingUnicastHostsPort` | x | - `networkHost` | x | - `networkBindHost` | x | - `networkPublishHost` | x | - `transportTcpPort` | x | - `transportAddresses` | - | x `clientTransportSniff` | - | x `clientTransportIgnoreClusterName` | - | x `clientTransportPingTimeout*` | - | x `clientTransportNodesSamplerInterval` | - | x `httpEnabled` | x | - `httpCORSEnabled` | x | - `httpCORSAllowOrigin` | x | - `httpCORSConfigurations` | x | - `additionalConfigurations` | x | x `additionalIndexConfigurations` | x | x `additionalTypeMappings` | x | x `overrideTypeMappings` | x | x `syncSearch` | x | -

* **Note:** Available in the Connector to Elasticsearch 6 only.

1 This is, of course, a nod to all those fans of Boaty Mcboatface.

88.2 Advanced Configuration of the Liferay Elasticsearch Connector (6.1)

The default configurations for Liferay's Elasticsearch adapter module are set in a Java class called `ElasticsearchConfiguration`.

While the Elasticsearch adapter has a lot of configuration options out of the box, you might find an Elasticsearch configuration you need that isn't provided by default. In this case, add the configuration options you need. If something is configurable for Elasticsearch, its configurable using the Elasticsearch adapter.

Adding Settings and Mappings to the Liferay Elasticsearch Adapter

The available configuration options are divided into two groups: the ones used most often by default and a catch-all for everything else. If the necessary setting isn't available by default, you can still configure it with the Liferay Elasticsearch adapter. Specify the settings you need by using one or more of the `additionalConfigurations`, `additionalIndexConfigurations`, or `additionalTypeMappings`, and `overrideTypeMappings` settings.

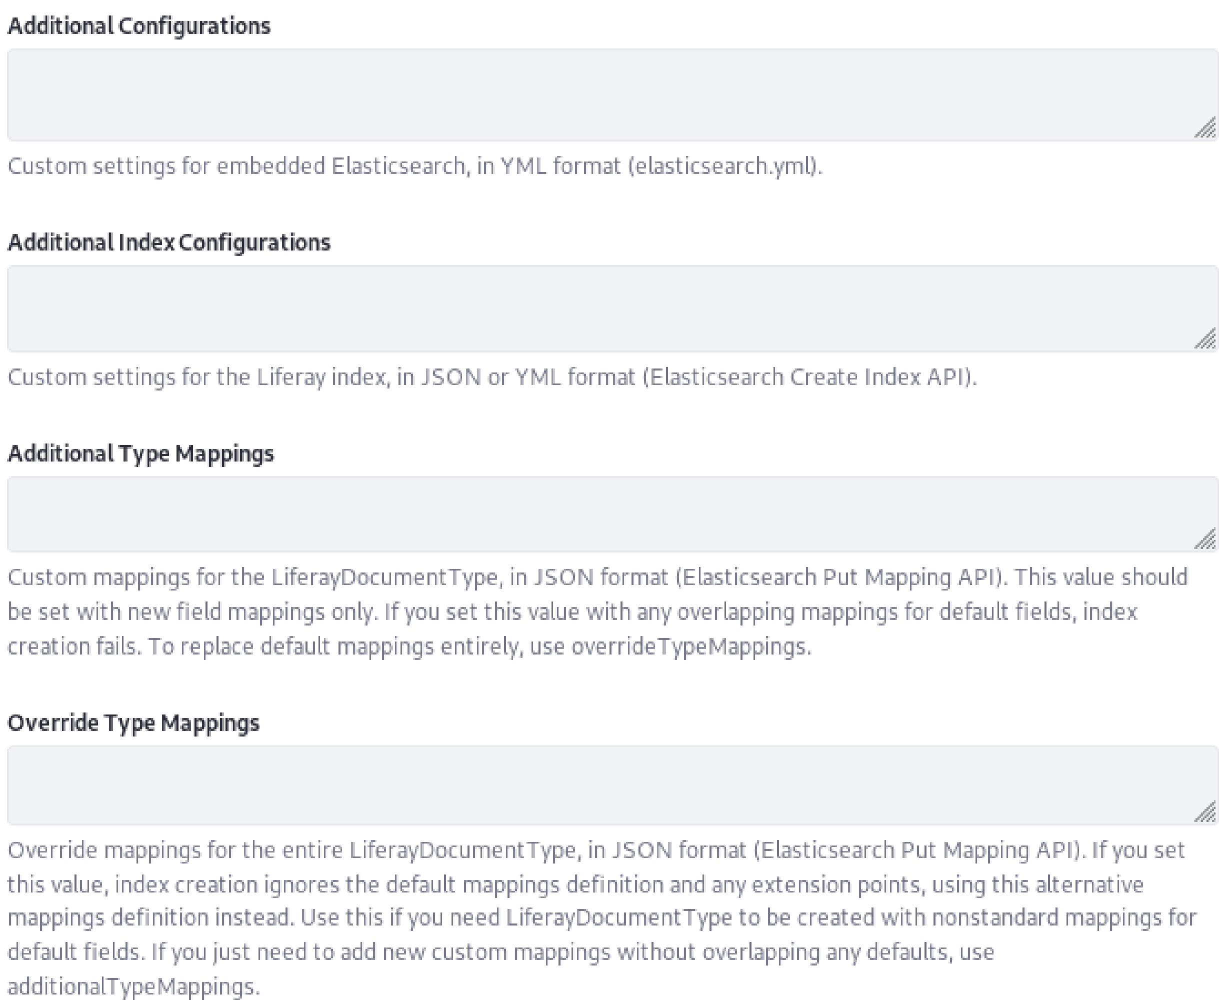

Figure 88.3: You can add Elasticsearch configurations to the ones currently available in System Settings.

`additionalConfigurations` defines extra settings (in YAML) for the embedded Elasticsearch. This is only useful for testing environments using the embedded Elasticsearch server. Any node settings normally set in `elasticsearch.yml` can be declared here. See the Elasticsearch documentation for a description of all possible node settings.

`additionalIndexConfigurations` defines extra settings (in JSON or YAML) that are applied to the Liferay DXP index when it's created. For example, you can create custom analyzers and filters using this setting. For a complete list of available settings, see the Elasticsearch reference.

Here's an example that shows how to configure analysis that can be applied to a dynamic template (see below).

```
{
    "analysis": {
        "analyzer": {
            "kuromoji_liferay_custom": {
                "filter": [
                    "cjk_width",
                    "kuromoji_baseform",
                    "pos_filter"
                ],
                "tokenizer": "kuromoji_tokenizer"
            }
        },
        "filter": {
            "pos_filter": {
                "type": "kuromoji_part_of_speech"
            }
        }
    }
}
```

`additionalTypeMappings` defines extra field mappings for the `LiferayDocumentType` type definition. These are applied when the index is created. Add these field mappings in using JSON syntax. For more information see here and here. Use `additionalTypeMappings` for new field mappings, but do not try to override existing `properties` mappings. If any of the `properties` mappings set here overlap with existing mappings, index creation fails. Use `overrideTypeMappings` to replace the default `properties` mappings.

Here's an example of a dynamic template that uses the analysis configuration above to analyze all string fields that end with `_ja`.

```
{
    "LiferayDocumentType": {
        "dynamic_templates": [
            {
                "template_ja": {
                    "mapping": {
                        "analyzer": "kuromoji_liferay_custom",
                        "index": "analyzed",
                        "store": "true",
                        "term_vector": "with_positions_offsets",
                        "type": "string"
                    },
                    "match": "\\w+_ja\\b|\\w+_ja_[A-Z]{2}\\b",
                    "match_mapping_type": "string",
                    "match_pattern": "regex"
                }
            }
        ]
    }
}
```

The above code adds a new `template_ja` dynamic template. This overrides the existing dynamic template with the same name. As with dynamic templates, you can add sub-field mappings to Liferay DXP's type mapping. These are referred to as properties in Elasticsearch.

```
{
    "LiferayDocumentType": {
        "properties": {
            "fooName": {
                "index": "true",
                "store": "true",
                "type": "keyword"
            }
        }
    }
}
```

See here for more details on Elasticsearch's field datatypes.

The above example shows how a `fooName` field might be added to Liferay DXP's type mapping. Because `fooName` is not an existing property in the mapping, it works fine. If you try to override an existing property mapping, index creation fails. Instead use the `overrideTypeMappings` setting to override `properties` in the mapping.

Note: To see that your additional mappings have been added to the `LiferayDocumentType`, use `curl` to access this URL after saving your additions and re-indexing:

```
curl http://[HOST]:[ES_PORT]/liferay-[COMPANY_ID]/_mapping/LiferayDocumentType?pretty
```

Here's what it would look like for an Elasticsearch instance running on `localhost:9200`, with a Liferay DXP Company ID of `20116`:

```
curl http://localhost:9200/liferay-20116/_mapping/LiferayDocumentType?pretty
```

In the above URL, `liferay-20116`is the index name. Including it indicates that you want to see the mappings that were used to create the index with that name.

Use `overrideTypeMappings` to override Liferay DXP's default type mappings. This is an advanced feature that should be used only if strictly necessary. If you set this value, the default mappings used to define the Liferay Document Type in Liferay DXP source code (for example, `liferay-type-mappings.json`) are ignored entirely, so include the whole mappings definition in this property, not just the segment you're modifying. To make a modification, find the entire list of the current mappings being used to create the index by navigating to the URL

```
http://[HOST]:[ES_PORT]/liferay-[COMPANY_ID]/_mapping/LiferayDocumentType?pretty
```

Copy the contents in as the value of this property (either into System Settings or your OSGi configuration file). Leave the opening curly brace `{`, but delete lines 2-4 entirely:

```
"liferay-[COMPANY_ID]": {
    "mappings" : {
        "LiferayDocumentType" : {
```

Then, from the end of the mappings, delete the concluding three curly braces.

```
        }
    }
}
```

Now modify whatever mappings you'd like. The changes take effect once you save the changes and trigger a re-index from Server Administration. If you need to add new custom mappings without overriding any defaults, use `additionalTypeMappings` instead.

Multi-line YAML Configurations

If you configure the settings from the last section using an OSGi configuration file, you might find yourself needing to write YAML snippets that span multiple lines. The syntax for that is straightforward and just requires appending each line with `\n\`, like this:

```
additionalConfigurations=\
                    cluster.routing.allocation.disk.threshold_enabled: false\n\
                    cluster.service.slow_task_logging_threshold: 600s\n\
                    index.indexing.slowlog.threshold.index.warn: 600s\n\
                    index.search.slowlog.threshold.fetch.warn: 600s\n\
                    index.search.slowlog.threshold.query.warn: 600s\n\
                    monitor.jvm.gc.old.warn: 600s\n\
                    monitor.jvm.gc.young.warn: 600s
```

From simple configurations to overriding existing type mappings, Elasticsearch and Liferay's connector to Elasticsearch are configurable.

Chapter 89

Installing Elasticsearch 7

Elasticsearrch 7 can be installed with Liferay DXP 7.1 (DXP-only). This article presents a basic setup. For additional details on the general Elasticsearch Installation procedure, refer to Installing Elasticsearch.

Once installed refere to the configuration guide.

Note: Elasticsearch 6.x and 7.x are supported for 7.0. Refer to the Search Engine Compatibility Matrix for details.

To install Elasticsearch 7 on a new 7.0 installation:

1. Install and configure Elasticsearch 7.x.
2. Install the Liferay Connector to Elasticsearch 7.
3. Disable the default, bundled Liferay Connector to Elasticsearch 6.
4. Configure Liferay DXP to connect to your Elasticsearch 7.x cluster.
5. Restart Liferay DXP and reindex your search indexes.

89.1 Install Elasticsearch

To install Elasticsearch,

1. Download an Elasticsearch archive (not the OSS version) from Elastic's website.

   Download the latest Elasticsearch archive compatible with your Liferay version.

2. Extract the archive contents to a local folder where you want to run Elasticsearch. This folder is your *Elasticsearch Home*.

3. Install the required Elasticsearch plugins by running these commands in your `[Elasticsearch Home]/bin` folder:

```
./elasticsearch-plugin install analysis-icu
```

```
./elasticsearch-plugin install analysis-kuromoji
```

```
./elasticsearch-plugin install analysis-smartcn
```

```
./elasticsearch-plugin install analysis-stempel
```

Each Elasticsearch server is configured by its `[Elasticsearch Home]/config/elasticsearch.yml` file.

Here's an example `elasticsearch.yml` configuration for a single-node cluster:

```
cluster.name: LiferayElasticsearchCluster

discovery.type: single-node
network.host: es-network
node.name: es-node1
transport.port: 9300

# Include additional settings (e.g., security settings)
```

This cluster is named `LiferayElasticsearchCluster` and has one node named `es-node1`.

If you are not configuring hosts for a production mode setup, use `localhost` as the network host value. Elasticsearch can bind to loopback addresses for HTTP and Transport communication. If you use a loopback network address and single node discovery, this means the Elasticsearch server is running in *development mode*.

Start Elasticsearch.

```
./bin/elasticsearch
```

89.2 Install the Liferay Connector to Elasticsearch 7

Liferay requires installation of a connector application that can enable communication with the corresponding Elasticsearch version:

1. Download the appropriate connector for your version of Liferay and Elasticsearch.
2. Stop Liferay and place the LPKG file into `Liferay Home/osgi/marketplace/`.

89.3 Disable the default Liferay Connector to Elasticsearch 6

Whether you're installing Elasticsearch 7 for a new 7.0 installation or upgrading an existing stack to Elasticsearch 6, you must disable Liferay's Elasticsearch 6 search modules bundled with the Liferay DXP by default.

1. Create a configuration file named

   com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.config

2. Give it these contents:

```
blacklistBundleSymbolicNames=[ \
    "com.liferay.portal.search.elasticsearch6.api", \
    "com.liferay.portal.search.elasticsearch6.impl", \
    "com.liferay.portal.search.elasticsearch6.spi", \
    "com.liferay.portal.search.elasticsearch6.xpack.security.impl", \
    "Liferay Connector to X-Pack Security [Elastic Stack 6.x] - Impl", \
    "Liferay Enterprise Search Security  - Impl" \
]
```

3. Place the file in `Liferay Home/osgi/configs`.

Note: After disabling the Elasticsearch 6 connection, you'll see exceptions in the Liferay log, because there's now no search engine connection. This is fixed by the next step.

89.4 Configure Liferay

Create a configuration file named

```
com.liferay.portal.search.elasticsearch7.configuration.ElasticsearchConfiguration.config
```

Give it these contents:

```
operationMode="REMOTE"
transportAddresses=["localhost:9300"]
```

This simple configuration leverages many default settings, but the configuration is customizable. See the Elasticsearch documentation to learn about the available transport settings.

Learn more about the Liferay connector's corresponding configurations in the configuration guide.

89.5 Restart and Re-Index

If both Liferay and Elasticsearch are running, once the connection is configured the Elasticsearch log shows that the Liferay indexes are created:

```
[2021-05-05T13:16:29,016][INFO ][o.e.c.m.MetadataCreateIndexService] [es-node1] [liferay-0] creating index, cause [api], templates [], shards [1]/[1
[2021-05-05T13:16:29,032][INFO ][o.e.c.r.a.AllocationService] [es-node1] updating number_of_replicas to [0] for indices [liferay-
0]
[2021-05-05T13:16:29,300][INFO ][o.e.c.r.a.AllocationService] [es-node1] Cluster health status changed from [YELLOW] to [GREEN] (reason: [shards sta
0][0]]]).
[2021-05-05T13:16:29,392][INFO ][o.e.c.m.MetadataMappingService] [es-node1] [liferay-0/rkLFWo4gT9mkL9Y2wrKkpA] update_mapping [LiferayDocumentType]
[2021-05-05T13:16:29,468][INFO ][o.e.c.m.MetadataMappingService] [es-node1] [liferay-0/rkLFWo4gT9mkL9Y2wrKkpA] update_mapping [LiferayDocumentType]
[2021-05-05T13:16:29,534][INFO ][o.e.c.m.MetadataCreateIndexService] [es-node1] [liferay-20099] creating index, cause [api], templates [], shards [1
[2021-05-05T13:16:29,535][INFO ][o.e.c.r.a.AllocationService] [es-node1] updating number_of_replicas to [0] for indices [liferay-
20099]
[2021-05-05T13:16:29,662][INFO ][o.e.c.r.a.AllocationService] [es-node1] Cluster health status changed from [YELLOW] to [GREEN] (reason: [shards sta
20099][0]]]).
[2021-05-05T13:16:29,722][INFO ][o.e.c.m.MetadataMappingService] [es-node1] [liferay-20099/214vtbkjTkmW_VCR_xOqyQ] update_mapping [LiferayDocumentTy
[2021-05-05T13:16:29,785][INFO ][o.e.c.m.MetadataMappingService] [es-node1] [liferay-20099/214vtbkjTkmW_VCR_xOqyQ] update_mapping [LiferayDocumentTy
```

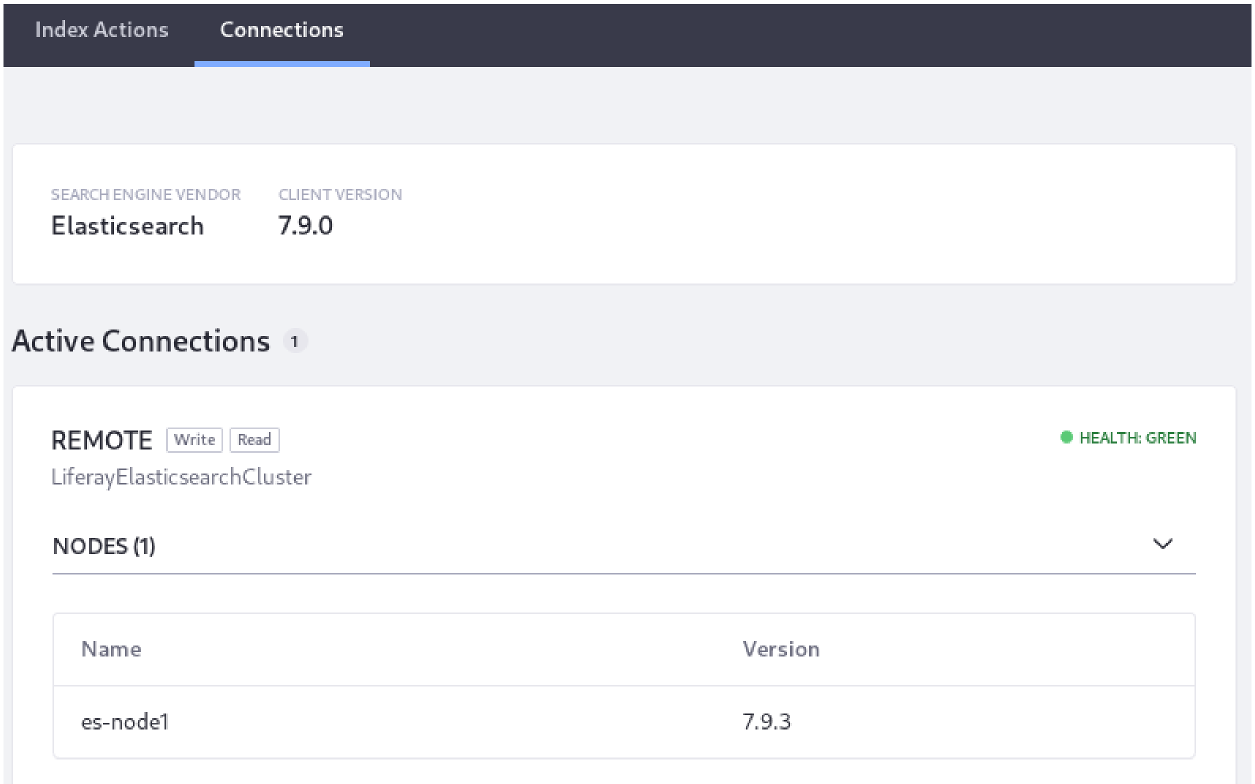

Figure 89.1: An active connection is displayed in the Search administrative panel.

In the Control Panel, navigate to Configuration → Search and verify the Elasticsearch connection is active.

Important: Re-index your search indexes and spell check indexes. Invoke both of these actions in the Index Actions tab of Control Panel → Configuration → Search.

89.6 Upgrading to Elasticsearch 7

Elasticsearch 7 is supported for Liferay DXP 7.1. If you're upgrading Liferay DXP and still running Elasticsearch 6, consider upgrading your Elasticsearch servers too. If you're setting up a new system and not already running a remote Elasticsearch 6 server, follow the installation guide to install Elasticsearch 7 and the configuration guide to configure the Elasticsearch adapter.

Before Proceeding, back up your existing data before upgrading Elasticsearch. If something goes wrong during or after the upgrade, roll back to the previous version using the uncorrupted index snapshots. See here for more information.

Here, you'll learn to upgrade an existing Elasticsearch 6 server (or cluster) to Elasticsearch 7:

1. Install and configure Elasticsearch 7.

2. In 7.0, security is now provided out of the box. If you're using X-Pack security, enable it (it's disabled by default):

   ```
   xpack.security.enabled: true
   ```

3. Blacklist the bundled Liferay Connector to Elasticsearch 6.

4. Install and configure the Liferay Connector to Elasticsearch 7.

5. Re-index all search and spell check indexes.

Learn about configuring Elasticsearch here.

Blacklisting Elasticsearch 6

To blacklist Elasticsearch 6,

1. Create a configuration file named

   ```
   com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.config
   ```

2. Give it these contents:

   ```
   blacklistBundleSymbolicNames=[ \
       "com.liferay.portal.search.elasticsearch6.api", \
       "com.liferay.portal.search.elasticsearch6.impl", \
       "com.liferay.portal.search.elasticsearch6.spi", \
       "com.liferay.portal.search.elasticsearch6.xpack.security.impl", \
       "Liferay Connector to X-Pack Security [Elastic Stack 6.x] - Impl", \
       "Liferay Enterprise Search Security  - Impl" \
   ]
   ```

Re-index

Once the Elasticsearch adapter is installed and talking to the Elasticsearch cluster, navigate to *Control Panel* → *Configuration* → *Search*, and click *Execute* for the *Reindex all search indexes* entry.

You must also re-index the spell check indexes.

Reverting to Elasticsearch 6

Stuff happens. If that stuff involves an unrecoverable failure during the upgrade to Elasticsearch 7, roll back to Elasticsearch 6 and regroup.

Since your Elasticsearch 6 and 7 are currently two separate installations, this procedure takes only a few steps:

1. Stop the Liferay Connector to Elasticsearch 6.

2. Stop Elasticsearch 7 and make sure that the Elasticsearch 6 `elasticsearch.yml` and the connector app are configured to use the same port (9200 by default).

3. Start the Elasticsearch server, and then restart the Liferay Connector to Elasticsearch 6.

Chapter 90

Installing Liferay Enterprise Search

A Liferay Enterprise Search (LES) subscription provides access to additional search features beyond what's available out of the box with your Liferay DXP subscription. To see a detailed description of the services and features included with LES on Liferay DXP, refer to the official description of LES in the Liferay DXP Components resource.

Note: To install Cross-Cluster Replication with 7.0, follow the instructions found on Liferay Learn. The procedure is nearly identical between Liferay 7.1 and 7.2, with any differences noted inline in the Liferay Learn articles.

LES customers receive a platinum Elasticsearch license from Liferay. There can be a delay between your subscription and receipt of the license, but you can enable a 30-day trial to work with in the meantime.

Detailed installation and usage instructions are available in the article for each LES feature, including

- [Liferay Learn] Cross-Cluster Replication
- Monitoring Elasticsearch
- Securing Elasticsearch [Free without LES for Liferay DXP with Elasticsearch 7]

Always check the LES compatibility matrix for compatibility information.

Here's an overview of using LES with Liferay DXP:

1. Get an Enterprise Search subscription.
2. You'll receive a license for the Elasticsearch features. Install it on your Elasticsearch servers.
3. Download and install the Liferay Enterprise Search apps you purchased.
4. Configure the connectors with the proper credentials and encryption information.
5. Restart Elasticsearch.

Chapter 91

Installing Liferay Enterprise Search Security

91.1 Enabling X-Pack Security

Elasticsearch 6.x: If you're using Elasticsearch 6, a Liferay Enterprise Search (LES) subscription is required to use X-Pack Security in your Liferay DXP-Elasticsearch stack. Starting with the *Liferay Connector to Elasticsearch 7* (available for Liferay DXP 7.1 on Liferay Marketplace), X-Pack security is included by default. Liferay Enterprise Search Monitoring still requires LES.

The first thing to do is enable X-Pack security. Add this setting in `elasticsearch.yml`:

```
xpack.security.enabled: true
```

Now you can set up X-Pack users.

91.2 Setting Up X-Pack Users in Elasticsearch

In a system using Security and Monitoring, Kibana and Liferay must authenticate as clients to Elasticsearch, using the credentials configured for these built-in X-Pack users:

- `kibana_system`
- `elastic`

Set the passwords for all X-Pack's built-in users. The `setup-passwords` command is the simplest method to set the built-in users' first-use passwords for the first time. To update a password subsequently, use Kibana's UI or the Change Password API.

The `interactive` argument sets the passwords for all built-in users. The configuration shown in these articles assumes you set all passwords to *liferay*. Of course, that's not recommended for production systems.

```
./bin/elasticsearch-setup-passwords interactive
```

Elastic's setup-passwords command documentation describes additional options.
Since you're securing Elasticsearch, remember the `elastic` user's password.
Enable transport layer security on each node.

91.3 Encrypting Communications in Elasticsearch

The following instructions for enabling SSL/TLS use `liferay` as the password whenever one is needed. Use your own passwords for your installation.

Important: Elasticsearch and Liferay DXP must share the keys and certificates used to configure SSL. Copy them between servers and point to the local copy in the corresponding configuration files.

Generate Node Certificates

Generate a certificate for each node, or generate a certificate to be used on all nodes and clients—like Liferay. Alternatively, use your certificate authority to obtain node certificates.

1. Generate X-Pack certificate authority using the X-Pack's `certutil` command:

   ```
   ./bin/elasticsearch-certutil ca --ca-dn CN=elastic-ca
   ```

   This generates a file called `elastic-stack-ca.p12`.

   To generate the certificate authority cert and private key in PEM format,

   ```
   ./bin/elasticsearch-certutil ca --pem --ca-dn CN=elastic-ca
   ```

2. Move the certificate authority file(s) file to the `[Elasticsearch Home]/config/certs` folder.

3. Generate X.509 certificates and private keys using the CA you created:

   To generate certificates and keys in `PKCS#12` format, run

   ```
   ./bin/elasticsearch-certutil cert --ca config/certs/elastic-stack-ca.p12 --ca-pass liferay --dns localhost --ip 127.0.0.1 --
   name elastic-nodes
   ```

   To generate certificates and keys in `PEM` format, run

   ```
   ./bin/elasticsearch-certutil cert --pem --ca-cert config/certs/ca.crt --ca-key config/certs/ca.key --dns localhost --
   ip 127.0.0.1 --name elastic-nodes
   ```

   To generate `PEM` format node certificates and keys from the `PKSC#12` certificate authority, run

   ```
   ./bin/elasticsearch-certutil cert --pem --ca config/certs/elastic-stack-ca.p12 --ca-pass liferay --dns localhost --
   ip 127.0.0.1 --name elastic-nodes
   ```

To generate a certificate that works with multiple hosts (for example to use the same certificate on all Elasticsearch and Liferay servers), use comma-separated lists when listing the DNS names and IP addresses:

```
./bin/elasticsearch-certutil cert --ca config/certs/elastic-stack-ca.p12 --ca-pass liferay --dns localhost,example.com,es-node1,es-node2,es-node3 --ip 127.0.0.1,<IPv4-address-2>,<IPv4-address-3>,<IPv4-address-4>
```

The `elasticsearch-certutil cert` command generates another file called `elastic-nodes.p12` (feel free to name it differently).

```
**Note:** The `certutil` command defaults to using the PKCS#12 format
for certificate generation, which works with your Elastic Stack 7.x. Kibana
6.x does not work with PKCS#12 certificates, so the `--pem` option
(generates the certificate in PEM format) is important if you're using
Liferay 7.1 and Kibana 6.x with *Liferay Enterprise Search Monitoring*. The
PEM command  for each case generates two ZIP files: `ca.crt` and
`ca.key`, `elastic-nodes.crt` and `elastic-nodes.key`. Unzip the
archive contents in `[Elasticsearch Home]/config/certs`.
```

---

1. Move `elastic-nodes.p12` to the `[Elasticsearch Home]/config/certs` folder.

   Checkpoint: You now have the following files in your `[Elasticsearch Home]/config/certs` folder:

   ```
   elastic-nodes.p12
   elastic-stack-ca.p12
   ```

   or

   ```
   ca.crt
   ca.key
   elastic-nodes.crt
   elastic-nodes.key
   ```

2. Copy the files to the same folder on each Elasticsearch node, and to an appropriate location on each Liferay server node.

The certificates and keys are ready to use in your Elasticsearch configuration.

Enable TLS in Elasticsearch

Enable TLS (6.x, 7.x) on each node via its `[Elasticsearch Home]/config/elasticsearch.yml` file.

1. Enable transport layer TLS with these settings in `elasticsearch.yml` for inter-node communication:

   ```
   xpack.security.transport.ssl.enabled: true
   ```

2. Add the certificate, key and certificate authority paths to each node's `elasticsearch.yml`:

```
# PKCS#12
xpack.security.transport.ssl.keystore.path: certs/elastic-nodes.p12
xpack.security.transport.ssl.keystore.password: liferay
xpack.security.transport.ssl.truststore.path: certs/elastic-nodes.p12
xpack.security.transport.ssl.truststore.password: liferay
xpack.security.transport.ssl.verification_mode: certificate
```

The example paths above assume you added the certificates to `[Elasticsearch Home]/config/certs`.

3. Enable TLS on the HTTP layer to encrypt client communication:

```
xpack.security.http.ssl.enabled: true
```

4. Configure the certificate, key, and certificate authority paths to each node's `elasticsearch.yml`:

```
# PKCS#12
xpack.security.http.ssl.keystore.path: certs/elastic-nodes.p12
xpack.security.http.ssl.keystore.password: liferay
xpack.security.http.ssl.truststore.path: certs/elastic-nodes.p12
xpack.security.http.ssl.truststore.password: liferay
```

PEM: To configure Elasticsearch using `PEM` format certificates, refer to the examples below.

Once TLS is enabled, configure X-Pack Security in Liferay DXP.

91.4 Configuring Security in Liferay's Elasticsearch Connector

---

Elasticsearch 6.x: If you have a Liferay Enterprise Search subscription, download the Liferay Enterprise Search Security app. Install the LPKG file by copying it into the `Liferay Home/deploy` folder.

Elasticsearch 7.x: The Liferay Connector to Elasticsearch 7 (DXP-only on Liferay 7.1) includes security integration by default; installing the LES Security app is unnecessary.

To configure security, navigate to *Control Panel* → *Configuration* → *System Settings*. Find the *Search* category and click on the *X-Pack Security* entry. You can enter the property values here, but it's more common to use a configuration file deployed to `[Liferay Home]/osgi/configs`.

On Elasticsearch 7.x, for the security adapter, create a file called

```
com.liferay.portal.search.elasticsearch7.configuration.XPackSecurityConfiguration.config
```

For Elasticsearch 6.x installations the file is named

```
com.liferay.portal.search.elasticsearch6.xpack.security.internal.configuration.XPackSecurityConfiguration.config
```

The exact contents of the file depend on your X-Pack setup in Elasticsearch. To configure the adapter according to the Elasticsearch setup documented here, populate the file like this (`PKCS#12`):

```
certificateFormat="PKCS#12"
sslKeystorePath="/PATH/TO/elastic-nodes.p12"
sslKeystorePassword="liferay"
sslTruststorePath="/PATH/TO/elastic-nodes.p12"
sslTruststorePassword="liferay"
requiresAuthentication=B"true"
username="elastic"
password="liferay"
transportSSLVerificationMode="certificate"
transportSSLEnabled=B"true"
```

Use settings like this if you're using `PEM` format certificates:

```
certificateFormat="PEM"
sslKeyPath="/PATH/TO/elastic-nodes.key"
sslCertificatePath="/PATH/TO/elastic-nodes.crt"
requiresAuthentication=B"true"
username="elastic"
password="liferay"
sslCertificateAuthoritiesPaths="/PATH/TO/ca.crt"
transportSSLVerificationMode="certificate"
transportSSLEnabled="true"
```

Note that the `password` should match what you set during the X-Pack password setup.

The certificate and key files referenced here are the same ones used on the Elasticsearch server. Copy them to the Liferay DXP server and update their paths in the configuration accordingly.

Enable authentication by setting `requiresAuthentication` to `true` and providing the credentials for the Elasticsearch user. For SSL, enable transport SSL, set the certificate verification mode and certificate format, and provide the path to the certificate, key, and certificate authority. Of course, the exact values depend on your X-Pack configuration.

Here's the complete list of security configuration options:

- `sslKeyPath`
- `sslCertificatePath`
- `sslCertificateAuthoritiesPaths`
- `certificateFormat`
- `requiresAuthentication`
- `username`
- `password`
- `transportSSLVerificationMode`
- `transportSSLEnabled`
- `sslKeystorePath`
- `sslKeystorePassword`
- `sslTruststorePath`
- `sslTruststorePassword`

When you're finished configuring X-Pack Security, restart Elasticsearch. These steps require a full cluster restart.

91.5 Example Elasticsearch Security Configurations

Here are the complete Elasticsearch 7.x configurations from this article (`elasticsearch.yml`; applies equally to Elasticsearch 6.5.x+):

Example Elasticsearch Security Configuration - PKCS#12

```
cluster.name: LiferayElasticsearchCluster

# X-Pack Security
xpack.security.enabled: true

## TLS/SSL settings for Transport layer (PKCS#12)
xpack.security.transport.ssl.enabled: true
xpack.security.transport.ssl.keystore.path: certs/elastic-nodes.p12
xpack.security.transport.ssl.keystore.password: liferay
xpack.security.transport.ssl.truststore.path: certs/elastic-nodes.p12
xpack.security.transport.ssl.truststore.password: liferay
xpack.security.transport.ssl.verification_mode: certificate

# TLS/SSL settings for HTTP layer (PKCS#12)
xpack.security.http.ssl.enabled: true
xpack.security.http.ssl.keystore.path: certs/elastic-nodes.p12
xpack.security.http.ssl.keystore.password: liferay
xpack.security.http.ssl.truststore.path: certs/elastic-nodes.p12
xpack.security.http.ssl.truststore.password: liferay

# Comment out when Kibana and Liferay's LES Monitoring are also configured
#xpack.monitoring.collection.enabled: true
```

91.6 Example Elasticsearch Security Configuration - PEM

```
cluster.name: LiferayElasticsearchCluster

# X-Pack Security
xpack.security.enabled: true

## TLS/SSL settings for Transport layer (PEM)
xpack.security.transport.ssl.enabled: true
xpack.security.transport.ssl.certificate: certs/elastic-nodes.crt
xpack.security.transport.ssl.certificate_authorities: ["certs/ca.crt"]
xpack.security.transport.ssl.key: certs/elastic-nodes.key
xpack.security.transport.ssl.key_passphrase: liferay
xpack.security.transport.ssl.verification_mode: certificate

# TLS/SSL settings for HTTP layer (PEM)
xpack.security.http.ssl.enabled: true
xpack.security.http.ssl.certificate: certs/elastic-nodes.crt
xpack.security.http.ssl.certificate_authorities: ["certs/ca.crt"]
xpack.security.http.ssl.key: certs/elastic-nodes.key
xpack.security.http.ssl.key_passphrase: liferay

# Use this setting when Kibana and Liferay's LES Monitoring are also configured
#xpack.monitoring.collection.enabled: true
```

Chapter 92

Installing Liferay Enterprise Search Monitoring

Monitor Elasticsearch with X-Pack Monitoring. First install X-Pack onto Elasticsearch and configure security if you're using X-Pack's security features. Then come back here for instructions on installing and configuring Kibana (the monitoring server) with X-Pack so that Elasticsearch (secured with X-Pack), Kibana (secured with X-Pack), and Liferay DXP can communicate effortlessly and securely. A Liferay Enterprise Search subscription is necessary for this integration. Contact Liferay's Sales department for more information.

1. Tell Elasticsearch to enable data collection.
2. Download and install Kibana.
3. Configure Kibana with the proper security settings.
4. Install the Liferay Enterprise Search Monitoring.
5. Configure the connector to communicate with Elasticsearch.

This document assumes you're enabling security *and* monitoring, though differences in the process are noted as appropriate.

Start by enabling data collection in Elasticsearch.

92.1 Enable Data Collection

Monitoring is enabled on Elasticsearch by default, but data collection isn't. Enable data collection by adding this line to `elasticsearch.yml`.

```
xpack.monitoring.collection.enabled: true
```

Now install Kibana.

92.2 Install Kibana

Make sure to install the correct version of Kibana. Check the Liferay Enterprise Search compatibility matrix for details.

1. Download Kibana and extract it. The root folder is referred to as *Kibana Home*.

2. Tell Kibana where to send monitoring data by setting Elasticsearch's URL in `kibana.yml`:

   ```
   elasticsearch.url: "http://localhost:9200"
   ```

   If SSL is enabled on Elasticsearch, this is an `https` URL.

3. If not using X-Pack security, start Kibana by opening a command prompt to Kibana Home and entering this command:

   ```
   ./bin/kibana
   ```

If you're using X-Pack's security features on the Elasticsearch server, there's additional configuration required before starting Kibana.

Configure Kibana with Authentication

If X-Pack requires authentication to access the Elasticsearch cluster, follow these steps or refer to Elastic's documentation.

1. Set the password for the built-in `kibana` user in `[Kibana Home]/config/kibana.yml`:

   ```
   elasticsearch.username: "kibana"
   elasticsearch.password: "liferay"
   ```

   Use your `kibana` user password from your X-Pack setup. Once Kibana is installed, you can change the built-in user passwords from the *Management* user interface.

2. If you're not encrypting communication with the Elasticsearch cluster, start Kibana from Kibana home.

   ```
   ./bin/kibana
   ```

3. Go to `localhost:5601` and make sure you can sign in as a user who has the `kibana_user` role or a superuser (like the `elastic` user).

Configuring Kibana with Encryption

Follow these steps to configure Kibana if X-Pack encrypts communication with the Elasticsearch cluster. Consult Elastic's guide for more information.

Add these settings to `kibana.yml`:

```
xpack.security.encryptionKey: "xsomethingxatxleastx32xcharactersx"
xpack.security.sessionTimeout: 600000

elasticsearch.ssl.verificationMode: certificate
elasticsearch.url: "https://localhost:9200"
elasticsearch.ssl.certificateAuthorities: [ "/path/to/ca.crt" ]

server.ssl.enabled: true
server.ssl.certificate: /path/to/[Elasticsearch Home]/config/localhost.crt
server.ssl.key: /path/to/[Elasticsearch Home]/config/localhost.key
```

For more information about monitoring and security best practices in a clustered environment, refer to Elastic's documentation.

After this step you can access Kibana at `https://localhost:5601` and sign in with a Kibana user. The last step is to connect Kibana to Liferay DXP.

92.3 Configuring the Liferay Enterprise Search Monitoring App

If you have a Liferay Enterprise Search (Premium or Standard) subscription, download the Liferay Enterprise Search Monitoring app. Install the LPKG file by copying it into the `Liferay Home/deploy` folder.

1. Once the connector is installed and Kibana and Elasticsearch are securely configured, create a configuration file named

   ```
   com.liferay.portal.search.elasticsearch6.xpack.monitoring.web.internal.configuration.XPackMonitoringConfiguration.config
   ```

2. Place these settings in the `.config` file:

   ```
   kibanaPassword="liferay"
   kibanaUserName="elastic"
   kibanaURL="http://localhost:5601"
   ```

   The values depend on your Kibana configuration. For example, use a secure URL such as `kibanaURL="https://localhost:5601"` if you're using X-Pack Security features.

   Alternatively, configure the monitoring adapter from System Settings. Navigate to *Control Panel* → *Configuration* → *System Settings* and find the X-Pack Monitoring entry in the Search category. All the configuration options for the monitoring connector appear there.

3. Deploy this configuration file to `Liferay Home/osgi/configs`, and your running instance applies the settings. There's no need to restart the server.

4. There are two more settings to add to Kibana itself. The first forbids Kibana from rewriting requests prefixed with `server.basePath`. The second sets Kibana's base path for the Monitoring portlet to act as a proxy for Kibana's monitoring UI. Add this to `kibana.yml`:

```
server.rewriteBasePath: false
server.basePath: "/o/portal-search-elasticsearch-xpack-monitoring/xpack-monitoring-proxy"
```

Note that once you set the `server.basePath`, you cannot access the Kibana UI through Kibana's URL (e.g., `https://localhost:5601`). All access to the Kibana UI is through the Monitoring portlet, which is only accessible to signed in Liferay DXP users. Navigate directly to the portlet using this URL:

http://localhost:8080/o/portal-search-elasticsearch-xpack-monitoring/xpack-monitoring-proxy/app/monitoring

5. Because you're using the Monitoring portlet in Liferay DXP as a proxy to Kibana's UI, if you are using X-Pack Security, you must configure the application server's startup JVM parameters to recognize a valid *truststore* and *password*.

   First, navigate to Elasticsearch Home and generate a PKSC#12 certificate from the CA you created when setting up X-Pack security:

```
./bin/elasticsearch-certutil cert --ca-cert /path/to/ca.crt --ca-key /path/to/ca.key --ip 127.0.0.1 --dns localhost --
name localhost --out /path/to/Elasticsearch_Home/config/localhost.p12
```

   Next use the `keytool` command to generate a truststore:

```
keytool -importkeystore -deststorepass liferay -destkeystore /path/to/truststore.jks -srckeystore /path/to/Elasticsearch_Home/config/localhos
srcstoretype PKCS12 -srcstorepass liferay
```

   Add the trustore path and password to your application server's startup JVM parameters. Here are example truststore and path parameters for appending to a Tomcat server's `CATALINA_OPTS`:

```
-Djavax.net.ssl.trustStore=/path/to/truststore.jks -Djavax.net.ssl.trustStorePassword=liferay
```

Restart Liferay DXP and Kibana.

92.4 Monitoring in Liferay DXP

Once Kibana and X-Pack are successfully installed and configured and all the servers are running, add the X-Pack Monitoring portlet to a page:

1. Open the *Add* menu on a page and choose *Widgets*
2. Search for *monitoring* and drag the *X-Pack Monitoring* widget from the Search category onto the page.

See the Elastic documentation for information on monitoring Elasticsearch.

Chapter 93

Installing X-Pack (6.1)

X-Pack is an Elasticsearch extension for securing and monitoring Elasticsearch clusters. If you use Elasticsearch, you should secure it with X-Pack. X-Pack's security features include authenticating access to the Elasticsearch cluster's data and encrypting Elasticsearch's internal and external communications. Most production systems requires these features. A Liferay Enterprise Search subscription gets you access to two X-Pack Connectors for Liferay DXP: monitoring and security. Contact Liferay's Sales department for more information.

Here's an overview of using X-Pack to secure the data indexed in Elasticsearch:

1. Get an Enterprise Search subscription.
2. Install X-Pack into Elasticsearch and configure it to require authentication and encryption.
3. Download and install the Liferay Enterprise Search Security.
4. Configure the secuirty app with the proper credentials and encryption information.
5. Restart Elasticsearch.

Important: These steps require a full cluster restart.

On completing these steps, your Elasticsearch installation communicates freely with Liferay DXP.

Elastic's documentation describes the X-Pack architecture and additional configuration options and features.

93.1 Installing X-Pack

Here are the X-Pack installation steps:

1. To install X-Pack and automatically grant it the required permissions (recommended), run this command on each Elasticsearch node:

```
bin/elasticsearch-plugin install x-pack --batch
```

The `--batch` option bypasses installation prompts for granting permissions to X-Pack.

The log output details the permissions granted and finishes with `Installed x-pack`:

```
-> Downloading x-pack from elastic
[=================================================] 100%
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@     WARNING: plugin requires additional permissions     @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
* java.io.FilePermission \\.\pipe\* read,write
* java.lang.RuntimePermission accessClassInPackage.com.sun.activation.registries
* java.lang.RuntimePermission getClassLoader
* java.lang.RuntimePermission setContextClassLoader
* java.lang.RuntimePermission setFactory
* java.net.SocketPermission * connect,accept,resolve
* java.security.SecurityPermission createPolicy.JavaPolicy
* java.security.SecurityPermission getPolicy
* java.security.SecurityPermission putProviderProperty.BC
* java.security.SecurityPermission setPolicy
* java.util.PropertyPermission * read,write
See http://docs.oracle.com/javase/8/docs/technotes/guides/security/permissions.html
for descriptions of what these permissions allow and the associated risks.
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@        WARNING: plugin forks a native controller        @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
This plugin launches a native controller that is not subject to the Java
security manager nor to system call filters.
Elasticsearch keystore is required by plugin [x-pack], creating...
-> Installed x-pack
```

See more about the permissions X-Pack needs here.

2. Make sure Elasticsearch does not allow automatic index creation by verifying this property in your `elasticsearch.yml` file:

```
action.auto_create_index: false
```

   This property is `true` by default. Elastic's documentation describes more automatic index creation details.

3. Shutdown Liferay DXP and restart Elasticsearch.

Now configure security and/or monitoring, depending on your needs.

93.2 Installing X-Pack Security (6.1)

Once X-Pack is installed, start securing Elasticsearch by configuring the built-in user passwords.

Setting Up X-Pack Users

In a system using X-Pack Security and X-Pack Monitoring, these built-in X-Pack users are important:

- `kibana`
- `elastic`

Set the passwords for all X-Pack's built-in users. The `setup-passwords` command is the simplest method to set the built-in users first-use passwords for the first time. To update a password subsequently, use Kibana's UI or the Change Password API.

The `interactive` argument sets the passwords for all built-in users. The configuration shown in these articles assumes you set all passwords to *liferay*. Of course, that's not recommended for production systems.

```
./bin/x-pack/setup-passwords interactive
```

Elastic's setup-passwords command documentation describes additional options.

Since you're securing Elasticsearch, remember the `elastic` user's password.

Enable transport layer security on each node.

Enabling Transport Layer Security

The following instructions for enabling TLS use `liferay` as the password whenever one is needed. Use your own passwords for your installation.

Generate Node Certificates

Generate a node certificate for each node. You can, of course, use a Certificate Authority to obtain node certificates.

1. Create a certificate authority, using X-Pack's `certutil` command:

   ```
   ./bin/x-pack/certutil ca --pem --ca-dn CN=localhost
   ```

   This generates a ZIP file. Unzip the contents somewhere safe.

2. Generate X.509 certificates and private keys using the CA from Step 1. For example:

   ```
   ./bin/x-pack/certutil cert --pem --ca-cert /path/to/ca.crt --ca-key /path/to/ca.key --dns localhost --ip 127.0.0.1 --
   name localhost
   ```

   This generates another ZIP file. Extract the contents somewhere in the `Elasticsearch Home/config` folder.

Note: The `certutil` command defaults to using the *PKSC#12* format for certificate generation. Since Kibana does not work with PKSC#12 certificates, the `--pem` option (generates the certificate in PEM format) is important if you're using monitoring.

Enable TLS

Enable TLS on each node via its `elasticsearch.yml`.

1. Add the certificate, key and certificate authority paths to each node's `elasticsearch.yml`:

   ```
   xpack.ssl.certificate: /path/to/[Elasticsearch Home]/config/localhost.crt
   xpack.ssl.key: /path/to/[Elasticsearch Home]/config/localhost.key
   xpack.ssl.certificate_authorities: ["/path/to/ca.crt"]
   ```

   The example paths above assume you added the certificate to `Elasticsearch Home/config/`.

2. Enable transport layer TLS with these settings in `elasticsearch.yml`:

   ```
   xpack.security.transport.ssl.enabled: true
   xpack.security.transport.ssl.verification_mode: certificate
   ```

3. Enable TLS on the HTTP layer to encrypt client communication:

   ```
   xpack.security.http.ssl.enabled: true
   ```

After X-Pack is installed and TLS is enabled, configure the security adapter in Liferay DXP.

Install and Configure the Liferay Enterprise Search Security app

If you have a Liferay Enterprise Search subscription, download the Liferay Enterprise Search Security app. Install the LPKG file by copying it into the `Liferay Home/deploy` folder. That's all there is to it.

To configure the security adapter, navigate to *Control Panel* → *Configuration* → *System Settings*. Find the *Search* category and click on the *X-Pack Security* entry. You can enter the property values here, but it's more common to use a configuration file deployed to `Liferay Home/osgi/configs`. For the security adapter, create a file called

```
com.liferay.portal.search.elasticsearch6.xpack.security.internal.configuration.XPackSecurityConfiguration.config
```

The exact contents of the file depend on your X-Pack setup. To configure the adapter according to the Elasticsearch setup documented here, populate the file like this:

```
sslKeyPath="/path/to/localhost.key"
sslCertificatePath="/path/to/localhost.crt"
certificateFormat="PEM"
requiresAuthentication="true"
username="elastic"
password="liferay"
sslCertificateAuthoritiesPaths="/path/to/ca.crt"
transportSSLVerificationMode="certificate"
transportSSLEnabled="true"
```

Note that the `password` should match what you set during the X-Pack password setup above.

The certificate and key files referenced here are the same ones used on the Elasticsearch server. Copy them to the Liferay DXP server and update their paths in the configuration accordingly.

Enable authentication by setting authentication to `required` and providing the credentials for the Elasticsearch user. For SSL, enable transport SSL, set the certificate verification mode and certificate format, and provide the path to the certificate, key, and certificate authority. Of course, the exact values depend on your X-Pack configuration.

Here's the complete list of configuration options for the security app:

- `sslKeyPath`
- `sslCertificatePath`
- `sslCertificateAuthoritiesPaths`
- `certificateFormat`
- `requiresAuthentication`
- `username`
- `password`
- `transportSSLVerificationMode`
- `transportSSLEnabled`
- `sslKeystorePath`
- `sslKeystorePassword`
- `sslTruststorePath`
- `sslTruststorePassword`

When you're finished configuring security, restart Elasticsearch. These steps require a full cluster restart.

93.3 Installing X-Pack Monitoring (6.1)

Monitor Elasticsearch with X-Pack Monitoring. First install X-Pack onto Elasticsearch and configure security if you're using X-Pack's security features. Then come back here for instructions on installing and configuring Kibana (the monitoring server) with X-Pack so that Elasticsearch (secured with X-Pack), Kibana (secured with X-Pack), and Liferay DXP can communicate effortlessly and securely. A Liferay Enterprise Search subscription is necessary for this integration. Contact Liferay's Sales department for more information.

1. Download and install Kibana.

2. Install X-Pack onto Kibana and configure Kibana with the proper security settings.

3. Download and install the Liferay Enterprise Search Monitoring app.

4. Configure the connector to communicate with Elasticsearch.

This document assumes you're enabling security *and* monitoring, though differences in the process are noted as appropriate.

For the X-Pack installation procedure, refer to the security article.

Start with the installation of Kibana.

Install Kibana

Make sure to install the correct version of Kibana. Check the Liferay Enterprise Search compatibility matrix for details.

1. Download Kibana and extract it. The root folder is referred to as *Kibana Home*.

2. Install X-Pack into Kibana:

```
./bin/kibana-plugin install x-pack
```

3. Tell Kibana where to send monitoring data by setting Elasticsearch's URL in `kibana.yml`:

   ```
   elasticsearch.url: "http://localhost:9200"
   ```

   If SSL is enabled on Elasticsearch, this is an `https` URL.

4. If not using X-Pack security, start Kibana by opening a command prompt to Kibana Home and entering this command:

   ```
   ./bin/kibana
   ```

If you're using X-Pack's security features on the Elasticsearch server, there's additional configuration required before starting Kibana.

Configure Kibana with Authentication

If X-Pack requires authentication to access the Elasticsearch cluster, follow these steps or refer to Elastic's documentation.

1. Set the password for the built-in `kibana` user in `Kibana Home/config/kibana.yml`:

   ```
   elasticsearch.username: "kibana"
   elasticsearch.password: "liferay"
   ```

   Use your `kibana` user password from your X-Pack setup. Once Kibana is installed, you can change the built-in user passwords from the *Management* user interface.

2. If you're not encrypting communication with the Elasticsearch cluster, start Kibana from Kibana home.

   ```
   ./bin/kibana
   ```

3. Go to `localhost:5601` and make sure you can sign in as a user who has the `kibana_user` role.

Configuring Kibana with Encryption

Follow these steps to configure Kibana if X-Pack encrypts communication with the Elasticsearch cluster. Consult Elastic's guide for more information.

Add these settings to `kibana.yml`:

```
xpack.security.encryptionKey: "xsomethingxatxleastx32xcharactersx"
xpack.security.sessionTimeout: 600000

elasticsearch.ssl.verificationMode: certificate
elasticsearch.url: "https://localhost:9200"
elasticsearch.ssl.certificateAuthorities: [ "/path/to/ca.crt" ]

server.ssl.enabled: true
server.ssl.certificate: /path/to/[Elasticsearch Home]/config/localhost.crt
server.ssl.key: /path/to/[Elasticsearch Home]/config/localhost.key
```

For more information about monitoring and security best practices in a clustered environment, refer to Elastic's documentation.

After this step you can access Kibana at `https://localhost:5601` and sign in with a Kibana user. The last step is to hook Kibana up with Liferay DXP.

Configuring the Liferay Enterprise Search Monitoring App

If you have a Liferay Enterprise Search (Premium or Standard) subscription, download the Liferay Enterprise Search Monitoring app. Install the LPKG file by copying it into the `Liferay Home/deploy` folder.

1. Once the connector is installed and Kibana and Elasticsearch are securely configured, create a configuration file named

   ```
   com.liferay.portal.search.elasticsearch6.xpack.monitoring.web.internal.configuration.XPackMonitoringConfiguration.config
   ```

2. Place these settings in the `.config` file:

   ```
   kibanaPassword="liferay"
   kibanaUserName="elastic"
   kibanaURL="http://localhost:5601"
   ```

   The values depend on your Kibana configuration. For example, use a secure URL such as `kibanaURL="https://localhost:5601"` if you're using X-Pack Security features.

   Alternatively, configure the monitoring adapter from System Settings. Navigate to *Control Panel* → *Configuration* → *System Settings* and find the X-Pack Monitoring entry in the Search category. All the configuration options for the monitoring connector appear there.

3. Deploy this configuration file to `Liferay Home/osgi/configs`, and your running instance applies the settings. There's no need to restart the server.

4. There's one more setting to add to Kibana itself. It sets Kibana's base path to let the Monitoring Portlet act as a proxy for Kibana's monitoring UI. Add this to `kibana.yml`:

   ```
   server.basePath: "/o/portal-search-elasticsearch-xpack-monitoring/xpack-monitoring-proxy"
   ```

   Note that once you set the `server.basePath`, you cannot access the Kibana UI through Kibana's URL (for example, `https://localhost:5601`). All access to the Kibana UI is through the Monitoring portlet, which is only accessible to logged in Liferay DXP users. Navigate directly to the portlet using this URL:

   http://localhost:8080/o/portal-search-elasticsearch-xpack-monitoring/xpack-monitoring-proxy/app/monitoring

5. Because you're using the Monitoring portlet in Liferay DXP as a proxy to Kibana's UI, if you are using X-Pack Security, you must configure the application server's startup JVM parameters to recognize a valid *truststore* and *password*.

   First, navigate to Elasticsearch Home and generate a PKSC#12 certificate from the CA you created when setting up X-Pack security:

   ```
   ./bin/x-pack/certutil cert --ca-cert /path/to/ca.crt --ca-key /path/to/ca.key --ip 127.0.0.1 --dns localhost --
   name localhost --out /path/to/Elasticsearch_Home/config/localhost.p12
   ```

   Next use the `keytool` command to generate a truststore:

```
keytool -importkeystore -deststorepass liferay -destkeystore /path/to/truststore.jks -srckeystore /path/to/Elasticsearch_Home/config/localhos
srcstoretype PKCS12 -srcstorepass liferay
```

Add the trustore path and password to your application server's startup JVM parameters. Here are example truststore and path parameters for appending to a Tomcat server's `CATALINA_OPTS`:

```
-Djavax.net.ssl.trustStore=/path/to/truststore.jks -Djavax.net.ssl.trustStorePassword=liferay
```

Restart Liferay DXP and Kibana.

Monitoring in Liferay DXP

Once Kibana and X-Pack are successfully installed and configured and all the servers are running, add the X-Pack Monitoring portlet to a page:

1. Open the *Add* menu on a page and choose *Widgets*
2. Search for *monitoring* and drag the *X-Pack Monitoring* widget from the Search category onto the page.

See the Elastic documentation for information on monitoring Elasticsearch and monitoring production systems.

Chapter 94

Installing Solr

Note: To install Solr 8, see the Solr installation guide on Liferay Learn

---

Solr is a popular enterprise search platform built on Apache Lucene. It's reliable, scalable, and fault tolerant. Read more about it here.

Although Elasticsearch is the default search engine that ships with Liferay DXP, it's perfectly valid to use Solr instead. In particular, if you've already been using Solr with a previous version of Liferay DXP, or your platform (for example, your OS or JVM) isn't supported by Elasticsearch, you might choose to use Solr to search and index your Liferay DXP data.

There are circumstances that force you to use Elasticsearch instead of Solr. Read here for more information.

7.0 supports Solr 7.5.x through the Liferay Connector to Solr 7 application.

94.1 Installing Solr: Basic Installation

Note: To install Solr 8, see the Solr installation guide on Liferay Learn

---

There are two ways to install the Liferay Connector to Solr 7:

1. Navigate to Liferay Marketplace and download the app that corresponds to your portal:

   ```
   **Liferay Portal CE:** [Liferay CE Connector to Solr 7](https://web.liferay.com/marketplace/-/mp/application/118014614)

   **Liferay DXP:** [Liferay Connector to Solr 7](https://web.liferay.com/marketplace/-/mp/application/117931595)

   Once the app LPKG is downloaded, copy it to
   `Liferay_Home/osgi/marketplace`.
   ```

2. In your running portal instance, navigate to *Control Panel* → *Apps* → *Store*. Sign in using your credentials, search for Solr Search Engine, and purchase (it's free) the Liferay Connector to Solr 7 entry.

This guide leads you through installing and configuring Solr. As you proceed, remember these terms:

Solr Home: The center of the Solr system (pun intended). This directory is `solr-[version]/server/solr`.

Liferay Home: The root folder of your Liferay DXP installation. It contains the `osgi`, `deploy`, `data`, and `license` folders, among others.

There are two installation steps:

1. Installing and configuring Solr 7.

2. Installing and configuring the Solr 7 connector for Liferay DXP.

Before configuring Liferay DXP for Solr, install and set up Solr.

Installing and Configuring Solr 7

To install and properly configure Solr for Liferay DXP:

1. Download Solr and unzip it.

2. Navigate to `solr-[version]/server/solr`. This is Solr Home.

3. Create a new folder called `liferay`.

4. In the `liferay` folder, create two new folders: `conf` and `data`.

5. Copy the contents of `Solr_Home/configsets/_default/conf` to `Solr_Home/liferay/conf`.

6. Open the Liferay Connector to Solr 7's LPKG file with an archive manager.

   Open the `com.liferay.portal.search.solr7.impl.jar` file, and extract

   ```
   META-INF/resources/solrconfig.xml
   ```

   and

   ```
   META-INF/resources/schema.xml
   ```

   to

   ```
   Solr_Home/liferay/conf
   ```

   This replaces the current `solrconfig.xml` and `schema.xml` files with ones that tell Solr how to index data coming from Liferay DXP.

7. Create a `core.properties` file in `Solr_Home/liferay` and add this configuration:

   ```
   config=solrconfig.xml
   dataDir=data
   name=liferay
   schema=schema.xml
   ```

8. Checkpoint: your `Solr_Home/liferay` folder now has this structure:

```
liferay
├── conf
│   ├── lang
│   │   ├── contractions_ca.txt
│   │   ├── ....txt
│   ├── managed-schema
│   ├── params.json
│   ├── protwords.txt
│   ├── schema.xml
│   ├── solrconfig.xml
│   ├── stopwords.txt
│   └── synonyms.txt
├── core.properties
└── data
```

9. Start the Solr server by entering

   ```
   ./bin/solr start -f
   ```

   from the top-level folder of your Solr installation (`solr-[version]`).

10. The Solr server listens on port `8983` by default. Navigate to `http://localhost:8983/solr/#/~cores` (assuming you're testing locally with `localhost` as your host), and confirm that the `liferay` core is available.

Solr is now installed. Next install and configure the Solr connector.

Installing and Configuring the Liferay Solr Adapter

Since Elasticsearch is the default search engine, the Elasticsearch connector is already installed and running. You must stop it before configuring the Solr connector.

Stop the Elasticsearch connector bundle using the App Manager, the Felix Gogo shell, or the bundle blacklist. If you're a Liferay DXP customer, use the blacklist feature as described below. The App Manager and Gogo shell rely on the `osgi/state` folder to "remember" the state of the bundle. If you delete this folder (recommended during patching) the Elasticsearch connector is reinstalled and started automatically.

Navigate to Control Panel → Apps → App Manager.

Once in the App Manager, search for *elasticsearch*. Find the Liferay Connector to Elasticsearch 6 module and click the Actions ((⋮)) button. Choose the Deactivate option. This leaves the bundle installed, but stops it in the OSGi runtime.

Alternatively, use the Felix Gogo shell to stop the Elasticsearch connector. Enter

```
lb elasticsearch
```

You'll see two active bundles for the Liferay Connector to Elasticsearch 6: an API and an IMPL bundle.

```
ID|State      |Level|Name
476|Active     |   10|Liferay CE Connector to Elasticsearch 6 - API (2.0.0)
478|Active     |   10|Liferay Portal Search Elasticsearch 6 API (2.0.6)
706|Active     |   10|Liferay CE Connector to Elasticsearch 6 - Impl (2.0.0)
707|Active     |   10|Liferay Portal Search Elasticsearch 6 Implementation (2.0.5)
```

Stop the API bundle by entering

```
stop [bundle ID]
```

In the example above, the `[bundle ID]` is `476`.

Liferay DXP: DXP customers should blacklist the Elasticsearch, Shield, and Marvel plugins.

1. Create a

   ```
   com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.config
   ```

   file with these contents:

   ```
   blacklistBundleSymbolicNames=["com.liferay.portal.search.elasticsearch6.api","com.liferay.portal.search.elasticsearch6.impl","Liferay Enterpr
   ```

   If the Liferay Enterprise Search LPKG files are installed, you must blacklist these too.

2. Place the file in `Liferay Home/osgi/configs`.

Install and configure the Solr connector:

1. Start Liferay DXP, then deploy the Solr connector by copying the LPKG you downloaded to `Liferay_Home/deploy`.

   You'll see a `STARTED` message in your Liferay DXP log once the Solr connector is installed. Here's what the log message looks like:

   ```
   2018-11-06 19:59:49.396 INFO  [pipe-start 943 944][BundleStartStopLogger:39] STARTED com.liferay.portal.search.solr7.api_1.0.0 [943]
   2018-11-06 19:59:49.490 INFO  [pipe-start 943 944][BundleStartStopLogger:39] STARTED com.liferay.portal.search.solr7.impl_1.0.0 [944]
   ```

2. To re-index against Solr, navigate to *Control Panel* → *Configuration* → *Search*, and click *Execute* next to the *Reindex all search indexes* option.

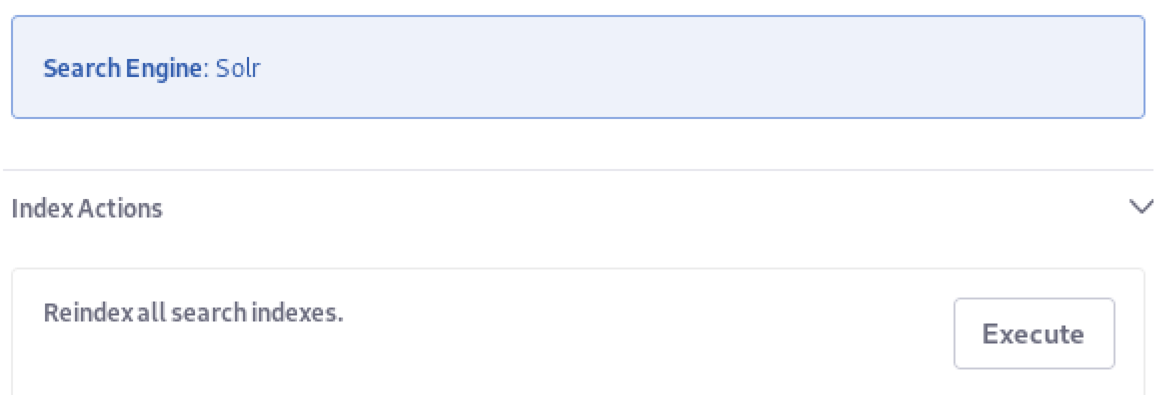

Figure 94.1: Once the Solr connector is installed, you can re-index your Liferay DXP data against your Solr server.

In production deployments, specify your edits to the Solr connector's default configurations using a configuration file deployed to the `Liferay_Home/osgi/configs` folder. Name the file

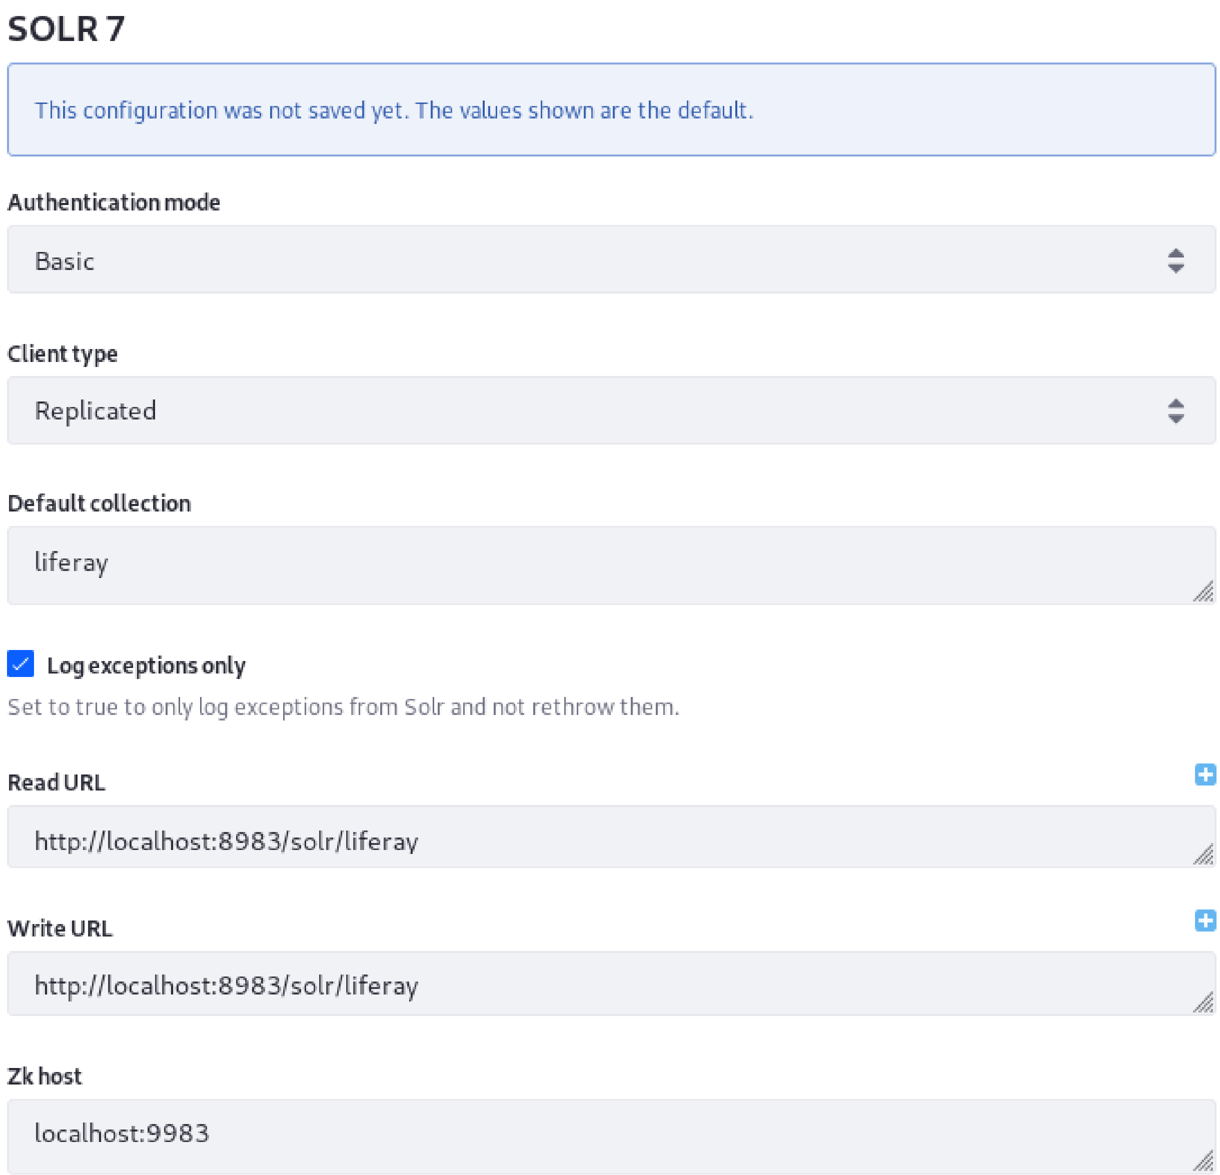

Figure 94.2: You can configure Solr from Liferay DXP's System Settings application. This is most useful during development and testing.

```
com.liferay.portal.search.solr7.configuration.SolrConfiguration.config
```

During testing and development, use the Solr 7 System Settings entry Control Panel → Configuration → System Settings for editing the default configurations.

The next article covers clustering Solr with SolrCloud.

94.2 Installing Solr: High Availability with SolrCloud

Note: To install Solr 8, see the Solr installation guide on Liferay Learn

---

Use SolrCloud if you need a cluster of Solr servers. Note that to use SolrCloud in production, you should set up an external ZooKeeper ensemble. ZooKeeper is a centralized coordination service for managing distributed systems like your SolrCloud cluster.

The steps included here should be considered the bare minimum of what must be done to configure SolrCloud with Liferay DXP. For example, these instructions cover configuring SolrCloud on a single machine, whereas a production environment would feature multiple physical or virtual machines. These instructions also assume you've followed the earlier section on *Installing and Configuring Solr 7*. Refer to the SolrCloud guide for more information.

1. Stop the Solr server if it's running.

2. Navigate to the `Solr_Home/configsets` folder and create a folder called

   ```
   liferay_configs
   ```

3. Copy the `conf` folder from `Solr_Home/liferay` to the `liferay_configs` folder you just created.

   The `configset/liferay_configs` folder contains the SolrCloud Liferay DXP collection configuration and is uploaded to ZooKeeper. By copying the `conf` folder from the `liferay` server configured earlier, you're using the `schema.xml` and `solrconfig.xml` files provided with the Liferay Solr Adapter.

4. Next launch an interactive SolrCloud session to configure your SolrCloud cluster. Use this command:

   ```
   ./bin/solr -e cloud
   ```

5. Complete the setup wizard. These steps demonstrate creating a two-node cluster:

   - Enter `2` for the number of nodes.
   - Specify ports `8983` and `7574` (the defaults). Both nodes are started with the start commands printed in the log:

     ```
     Starting up Solr on port 8983 using command:
     "bin/solr" start -cloud -p 8983 -s "example/cloud/node1/solr"
     ```

   - Name the collection *liferay*.
   - Split the collection into two shards.

- Specify two replicas per shard.
- When prompted to choose a configuration, enter *liferay_configs*. You should see a log message that concludes like this when the cluster has been started:

  ```
  SolrCloud example running, please visit http://localhost:8983/solr
  ```

Now you have a new collection called *liferay* in your local SolrCloud cluster. Verify its status by running the *status* command:

```
./bin/solr status
```

You'll see log output like this:

```
Found 2 Solr nodes:

Solr process 16989 running on port 8983
INFO  - 2018-08-06 13:54:17.665; org.apache.solr.util.configuration.SSLCredentialProviderFactory; Processing SSL Credential Provider chain: env;syspr
{
  "solr_home":"/home/russell/liferay-bundles/solr-7-dxp/solr-7.4.0/example/cloud/node1/solr",
  "version":"7.4.0 9060ac689c270b02143f375de0348b7f626adebc - jpountz - 2018-06-18 16:55:13",
  "startTime":"2018-08-06T17:52:01.519Z",
  "uptime":"0 days, 0 hours, 2 minutes, 16 seconds",
  "memory":"68.5 MB (%14) of 490.7 MB",
  "cloud":{
    "ZooKeeper":"localhost:9983",
    "liveNodes":"2",
    "collections":"1"}}


Solr process 17127 running on port 7574
INFO  - 2018-08-06 13:54:18.507; org.apache.solr.util.configuration.SSLCredentialProviderFactory; Processing SSL Credential Provider chain: env;syspr
{
  "solr_home":"/home/russell/liferay-bundles/solr-7-dxp/solr-7.4.0/example/cloud/node2/solr",
  "version":"7.4.0 9060ac689c270b02143f375de0348b7f626adebc - jpountz - 2018-06-18 16:55:13",
  "startTime":"2018-08-06T17:52:11.987Z",
  "uptime":"0 days, 0 hours, 2 minutes, 6 seconds",
  "memory":"56.4 MB (%11.5) of 490.7 MB",
  "cloud":{
    "ZooKeeper":"localhost:9983",
    "liveNodes":"2",
    "collections":"1"}}
```

To stop Solr while running in SolrCloud mode, use the *stop* command, like this:

```
bin/solr stop -all
```

Configure the Solr Adapter for SolrCloud

There's only one thing left to do: specify the client type as *CLOUD* in Liferay's Solr connector.

1. From System Settings or your OSGi configuration file, set the *Client Type* to *CLOUD*.

   ```
   clientType="CLOUD"
   ```

2. Start Liferay DXP if it's not running already.

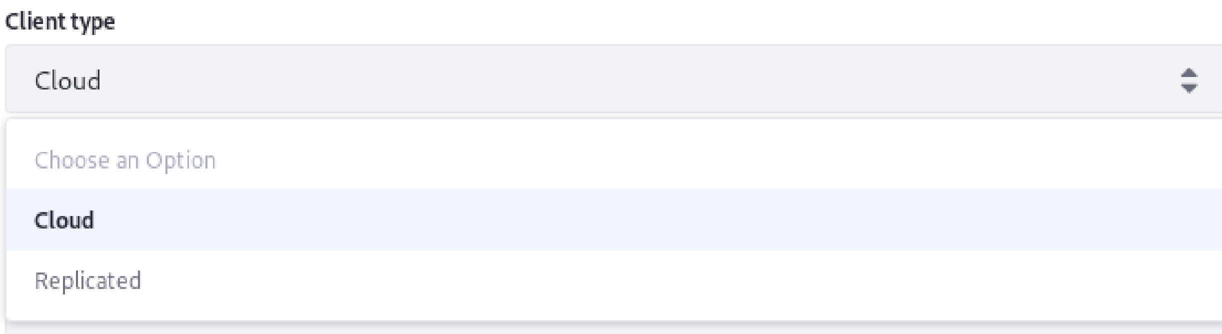

Figure 94.3: From the Solr 7 System Settings entry, set the *Client Type* to *Cloud*.

WC template | Template is deleted | The feed can exist without the template. It is not validated. | Renderer template is deleted | The feed can exist without the template. It is not validated. Page | Page's friendly URL is modified | The link to the page still works. | Portlet is removed from page | The feed can live without a portlet ID; the portlet is removed upon publication. Display page | WC article with display pages is configured | The page, Asset Publisher portlet, and content are published. | Display page is deleted | The page and its content are deleted. | WC article is deleted | The article deletion is published; the display page is not affected. | WC article with a display page is added but the page was not published earlier and it's not included in the publication | The page reference is validated and the publication fails when the page is not there. A message displays explaining this to the user.

Wiki

The following table describes how entities that are attached/related to Wikis are handled by the Staging framework.

Related entity	Action performed	How does Staging handle this?
Attachment link	Add a link to an attachment of the same page to a wiki page	The attachment is published and the link in the wiki points to the new attachment.
Add a link to another wiki page of the same node to a wiki page	The wiki page is published, but the link is not.	
Parent wiki page	Add a wiki page and a child of that page	Both pages are published.
Parent page is deleted	The parent page and all its children are deleted.	
Child page is deleted	The child page deletion is published.	
Parent page is updated	The updated page is published.	

Related entity	Action performed	How does Staging handle this?
Child page is updated	The updated page is published.	
Wiki redirect page	Page is moved to a new title	The pages are updated and published; the redirect works on the live site.
Intermediate page is deleted	The page deletion is published; the new page is not affected.	
Target page is deleted	Both the intermediate and new page are deleted and those deletions are published.	
Site link in content	Add a link to a Liferay site page to wiki	The new wiki content and the site page referenced in the link are published.
Wiki node	New node is added	The node is published.
Node is updated	The node is published.	
Set of pages is imported into existing node	The pages are published.	
Node is deleted	The node and all its dependent pages are deleted and those deletions are published.	
Wiki page	Page with no child/parent relationships is copied	The new page is published.
Page containing a child is moved, changing its title and parent	The page is published. After publication, the page is located under a new parent and has a new name. The child points to the new page.	
Page with no child/parent relationships is deleted	The page deletion is published.	
Attachment is removed from page	The page is published. The attachment is removed from the page on the live site.	
Category is removed from page	The page is published. The category is removed from the page on the live site.	
Page format is modified	The page is published.	

Related entity	Action performed	How does Staging handle this?
Image (not yet published) is added to page's content	The page and its image are published.	